



Europäisches Patentamt
European Patent Office
Office européen des brevets



(11) EP 0 953 899 A2

(12) EUROPEAN PATENT APPLICATION

(43) Date of publication:
03.11.1999 Bulletin 1999/44

(51) Int. Cl.⁶: G06F 9/44

(21) Application number: 99104072.6

(22) Date of filing: 17.03.1999

(84) Designated Contracting States:
AT BE CH CY DE DK ES FI FR GB GR IE IT LI LU
MC NL PT SE
Designated Extension States:
AL LT LV MK RO SI

- Saulpaugh, Thomas
San Jose, California 95120 (US)
- Bopardikar, Sunil K.
Sunnyvale, California 96087 (US)
- Li, Zi-Hui
Mountain View, California 94040 (US)

(30) Priority: 26.03.1998 US 48333

(71) Applicant:
Sun Microsystems, Inc.
Palo Alto, California 94043 (US)

(74) Representative:
Goddard, Heinz J., Dr.
FORRESTER & BOEHMERT
Franz-Joseph-Strasse 38
80801 München (DE)

(72) Inventors:
• Slaughter, Gregory K.
Palo Alto, California 94306 (US)

(54) Apparatus and method for object-oriented memory system

(57) A method and apparatus for providing memory functionality to object-oriented client software components of a computer system having a CPU using a first set of memory classes, each class of the first set being platform-independent; a second set of memory classes, each class of the second set being a subclass of a class

of the first set and being platform-dependent; and performing client component memory functions by accessing only objects of classes of the first set. Descriptions of the classes and methods are provided, enabling platform-independent device drivers to be implemented.




FIG. 1

Description

BACKGROUND OF THE INVENTION

Field of the Invention

[0001] The invention relates generally to computer memory and, more particularly, to memory methods and systems which are platform-independent.

Description of the Related Art

[0002] In order for computer systems to communicate with the user, the systems usually include a number of peripheral devices such as display screen, printers, and keyboards. Each of these peripheral devices requires a special software component, called a device driver, to provide for orderly exchange of data between the peripheral device and the rest of the computer system.

[0003] A company developing, for example, a new color printer would like to insure that the printer will be usable with computers having different hardware configurations ("platforms"), including systems based upon the most popular types of central processing units, such as the SPARC manufactured by Sun Microsystems, the PowerPC manufactured by the Motorola Corporation, and the Pentium manufactured by Intel Corporation. Currently, this requires that the printer manufacturer write a separate device driver for each platform, thereby driving up the development cost of the printer. Moreover, the need for custom device drivers for each platform also means that when a new platform is introduced, device drivers may not be immediately available for the most popular types of peripheral devices.

[0004] The need for custom device drivers for each platform is closely related to the varying memory characteristics of each platform. It would be desirable to provide memory systems and methods which would permit a single device driver to be written for a peripheral device, thereby allowing operation of the peripheral device on all platforms, including new platforms to be introduced in the future.

SUMMARY OF THE INVENTION

[0005] Features and advantages of the invention will be set forth in the description which follows, and in part will be apparent from the description, or may be learned by practice of the invention. The objectives and other advantages of the invention will be realized and attained by the apparatus, methods, and articles of manufacture particularly pointed out in the written description and claims hereof, as well as the appended drawings.

[0006] To achieve these and other advantages, and in accordance with the purpose of the invention as embodied and broadly described, the invention provides a method for providing memory functionality to object-oriented client software components of a computer system

having a CPU. The method comprising the steps of providing a first set of memory classes, each class of the first set being platform-independent; providing a second set of memory classes, each class of the second set being a subclass of a class of the first set and being platform-dependent; and performing client component memory functions by accessing only objects of classes of the first set.

[0007] In another aspect, the invention includes apparatus for performing memory functions in an object-oriented computer system, comprising a first set of memory classes, each class of the first set being platform-independent; a second set of memory classes, each class of the second set being a subclass of a class of the first set and being platform-dependent; a device driver accessing only objects of classes of the first set; and a bus manager accessing classes of the second set.

[0008] In yet another aspect, the invention includes a computer-readable medium containing instructions for providing memory functionality to object-oriented client software components of a computer system having a CPU. The instructions comprise a first set of memory classes, each class of the first set being platform-independent; and a second set of memory classes, each class of the second set being a subclass of a class of the first set and being platform-dependent.

[0009] It is to be understood that both the foregoing general description and the following detailed description are exemplary and explanatory and are intended to provide further explanation of the invention as claimed.

[0010] The accompanying drawings are included to provide a further understanding of the invention and are incorporated in and constitute a part of this specification, illustrate one/several embodiment(s) of the invention and, together with the description, serve to explain the principles of the invention.

BRIEF DESCRIPTION OF THE DRAWINGS

[0011] The accompanying drawings, which are incorporated in and constitute a part of this specification, illustrate embodiments of the invention and, together with the description, serve to explain the objects, advantages, and principles of the invention.

[0012] In the drawings:

Fig. 1 is a hardware block diagram of a computer system embodying the present invention.

Fig. 2 is a diagram of object-oriented software of the computer system of Fig. 1.

Fig. 3 is a diagram showing the class hierarchy of the memory system portion of the software shown in Fig. 2;

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENTS

[0013] Methods and apparatus consistent with the present invention are contained in an object-oriented operating system. The disclosed embodiment is implemented in the Java programming environment provided by Sun Microsystems, Inc. of Mountain View, CA. However, the invention is not so limited and may be incorporated into other computer systems, other operating systems, and other programming environments, as is understood by those skilled in the art. Java, Java-related trademarks, Sun, Sun Microsystems, and the Sun logo are either trademarks or registered trademarks of Sun Microsystems, Inc. of Mountain View, CA.

[0014] Reference will now be made in detail to an implementation consistent with the present invention, as illustrated in the drawings. Wherever possible, the same reference numbers will be used throughout the drawings and the following description to refer to the same or like parts.

[0015] Fig. 1 shows a computer system 10 which includes the present invention. Computer system 10 includes a central processing unit (CPU) 12 which may be, for example, a Sun SPARC, a Motorola Power PC or an Intel Pentium. Computer system 10 may represent a wide variety of computing devices. For example, system 10 may represent a standard personal computer, as widely used in homes and offices. Alternatively, system 10 may comprise a much more specialized "smart" system, such as a set-top box for use in receiving high definition television, or a multi-function cellular telephone.

[0016] CPU 12 is connected to memory 14 which may include various types of memory such as random access memory (RAM) and read only memory (ROM). CPU 12 is also connected to an expansion bus 16 which may be, for example, a PCI bus. Various types of input and output devices 18, 20, and 22 are connected to bus 16. For example, input device 18 may be a modem or network interface card connecting system 10 to a telephone line or local area network. Input device 20 may be, for example, a keyboard and output device 22 may be, for example, a printer. System 10 may optionally include a mass storage device 24, such as a hard disk drive connected by an I/O bus 26 which may be, for example, a SCSI bus. System 10 also includes a direct memory access (DMA) controller 23, under control of CPU 12, which can provide direct transfer of data between memory 14 and PCI bus 16.

[0017] Referring now to Fig. 2, there is shown a diagram of software stored in memory 14 of Fig. 1. Fig. 2 shows software logically arranged in a series of layers. An upper layer 30 is "platform-independent." The term "platform" generally refers to a CPU, physical memory, and permanently attached devices and buses. Thus, the software contained in platform-independent layer 32 is written in a manner that it may be used with any platform, based on any CPU, either existing or to be devel-

oped in the future, without modification. A second layer 34 is platform-dependent. Thus, software of layer 34 must be customized for the particular platform of computer system 10.

[0018] Platform-independent layer 32 includes an application program layer 36 which performs specific operations for the user, such as desktop publishing, managing a telephone call, or database management. Application layer 36 interfaces with a runtime system 38 which includes a component known as a "Java™ Virtual Machine" (JVM). The JVM is a software component which receives instructions in the form of machine-independent bytecodes produced by applications programs and interprets the instructions by dynamically converting and executing them. The JVM interfaces with a specific platform to execute desired functions. The disclosed embodiment employs a Java Virtual Machine, but other types of virtual machines are usable, as known to those skilled in the art. The operation of the JVM is well known to those skilled in the art and is discussed in, for example, the book Java! by Tim Ritchey, published by New Riders Publishing of Indianapolis, Indiana, and in The Java Virtual Machine Specification, by Lindham and Yellin, Addison-Wellesley (1997).

[0019] Runtime system 38 also includes a device interface portion 42 which supports operation of devices such as buses 16 and 26 and devices 18, 20, 22 (Fig. 1). In particular, device interface 42 includes device managers 46 and miscellaneous managers 48. Device interface 42 also includes device drivers 50 which are object-oriented programs written for each of the devices 18, 20, and 22. Note that device drivers 50 are included in platform-independent layer 32, and thus, once written, may be used to support devices 18, 20, and 22 on any platform both existing and to be developed in the future. Similarly, device interface 42 includes platform-independent bus managers 53 which are object-oriented programs written for buses, such as PCI bus 16 and SCSI bus 26 of Fig. 1. Device interface 42 also includes memory classes 52 which form part of the present invention and will be discussed below in greater detail.

[0020] Device interface 42 also includes additional components such as a system database 56 and a system loader 54 which support operation of computer system 10. Similarly, runtime layer 38 includes additional functions 58 supporting operations such as input/output, network operations, graphics, printing, multimedia, etc.

[0021] Platform-dependent layer 34 includes a platform interface 60, an OS native methods layer 61, a microkernel 62, and a boot interface 64. Platform interface 60 includes virtual machine system function library handlers 66 which may be written in the Java programming language to provide support for system function calls received from virtual machine 40. Platform interface 60 further includes interrupt classes 68 also written in the Java programming language which support inter-

rupt operations of computer system 10 and bus managers 44. Finally, platform interface 60 includes direct memory access (DMA) classes 70 and memory classes 72, each written in the Java programming language. DMA classes 70 and memory classes 72 are a feature of the present invention and will be discussed below in greater detail.

[0022] OS native methods layer 61 includes operating system methods of classes 68, 70, 72, and 44 which are written in a language specific to CPU 12 ("native" language) and which interface with lower-level operations of microkernel 62. These methods include interrupt native methods 78, DMA native methods 80, and memory native methods 82. Methods 80 and 82 will be discussed below in greater detail.

[0023] Microkernel 62 consists of native-language software components which support essential low-level hardware functions of CPU 12, such as resource allocation, interrupt process, security, etc. In particular, microkernel 62 includes a plurality of kernel functions 74 including thread management, exceptions, timing, physical memory management, hardware interrupt processing, platform control, process management, library management, I/O support, and monitor functions. Microkernel 62 further includes debug functions 76. Functions 74 and 76 may be performed by the Chorus microkernel commercially available from Sun Microsystems. However, other microkernels may be used, as is known to those skilled in the art.

[0024] The final component of platform-dependent layer 34 is boot interface 64. Boot interface 64 provides for loading and initialization of software into memory 14 (Fig. 1) when computer system 10 is initially powered up. Boot interface 64 may load software stored on, for example, a floppy disk, mass storage 24 (Fig. 1), or software received in the form of a computer data carrier wave over a network, via input device 18. Boot interface 64 forms no part of the present invention and will not be discussed in further detail.

[0025] Referring now to Fig. 3, there is shown a diagram illustrating the hierarchy of memory and DMA classes 52, 70, and 72. Memory and DMA Classes 52, 70, and 72 have the following design goals:

- One platform-independent class for client software components, such as drivers, to access (read or write) memory: Although there could be a variety of different types of platforms supported, a driver will always use the same platform-independent methods of these classes, and only these methods, to read and write memory. Developing device drivers written in the Java programming language, which are completely independent of the particular platform, is one of the main goals of the present invention.
- Methods for Bus Managers to Construct memory objects: the present invention encapsulates device platform specifics in a handful of bus managers

which construct memory objects for the vast array of device drivers to use to access memory. The Memory classes allow a bus manager to construct a generic memory object, which ultimately provides access to one of three (as specified by the bus manager) fundamental types of memory:

- Regular memory which is accessed via normal load and store operations from the CPU. This memory could be directly addressed by the CPU (i.e., physical memory), or could be indirectly accessed through an MMU (i.e. virtual memory).
- I/O memory which is accessed via normal load and store operations from the CPU, but which may also require other operations to be done to keep the I/O memory coherent with main memory. Again, this memory could be directly addressed by the CPU (i.e. physical memory), or could be indirectly accessed through an MMU (i.e. virtual memory).
- I/O port memory which is accessed via special I/O instructions (e.g. in and out instructions on Intel x 86-based platforms).
- Methods for the Bus Managers to set up DMA on behalf of client drivers.
- Platform independence of the Java programming language portions of the Memory and DMA classes themselves.
- Support for virtual memory and multiple address spaces. Furthermore, the virtual memory supported has to be sufficiently general to include platforms that have no MMU. In such cases, virtual address will be the same as physical address, and only one virtual address space will be provided by the platform. Device drivers will always think they are accessing virtual memory, even on systems with no MMU and only one virtual address space, even though they are accessing physical memory directly.
- A sufficiently rich memory model.
- A generic interface to a microkernel, via native methods, as shown in Fig. 2. The new Memory and DMA classes are designed to interface with a generic microkernel.

[0026] It is important to note in particular that the Memory and DMA classes do not provide yet another OS memory management subsystem implementation. Rather, the functionality for managing virtual memory, physical memory, I/O memory, and DMA is provided by the microkernel. The Memory and DMA classes do call upon such microkernel functionality via native methods, but these classes leverage that functionality, rather than reinventing it.

[0027] More detailed discussion of the classes shown in Fig. 3 is set forth below.

Memory Abstract Class

[0028] At the highest level, each of the Memory and DMA classes is eventually a sub-class of the *Memory* abstract class, which only has the notion of such general attributes as a base address, length, and constraints (of both access and allocation). Below that, memory is either ultimately accessible (perhaps not always directly, as in the case of physical memory) by a CPU, hence *MainMemory*, or it is accessed externally to the CPUs via DMA, hence *DMAMemory*.

MemoryDescriptor Class

[0029] Having a non-abstract and hence instantiable sub-class of *Memory* proves useful for representing fundamental characteristics of memory, such as base address and length, before an actual memory object has been constructed. Other than being instantiable, the *MemoryDescriptor* class does not provide any additional functionality beyond that provided by *Memory*. This is useful for bus managers to represent abstract memory before it is actually allocated.

MainMemory Abstract Class

[0030] The *MainMemory* abstract class only specifies abstract methods for managing caching, which is the only functionality common to its sub-classes. Eventually in the hierarchy (in *PhysicalMemory*, *VirtualMemory*, and *PortIOMemory*), these abstract cache management methods are implemented. It is important to note that cache management is highly platform-dependent, and thus each of these implementations ultimately relies on invoking the microkernel to flush caches, set the cache mode, etc. Therefore, the device driver has full access to complete cache functions using platform-independent methods of *MainMemory*.

MainMemory is either the abstract class, *AccessibleMemory*, which represents all memory that can be directly accessed by a CPU, or it is the *PhysicalMemory* class, which represents physical memory, which in general may not always be directly accessed by a CPU.

DMAMemory Class

[0031] The *DMAMemory* class provides methods for setting up DMA mappings to physical memory (ultimately done by the microkernel running below these classes), and for doing the corresponding unmapping. The decision was made to have the microkernel ultimately set up DMA mappings, since the full range of DMA hardware to be supported would not be something that it would be good to try to abstract at the *DMAMemory* class level, since doing so would be reinventing functionality that the microkernel must already support, and since it would probably not be possible to

develop a sufficiently general scheme that deals with all the conceivable flavors of DMA hardware.

[0032] Note in particular that the *DMAMemory* object is not concerned with whether the DMA addresses that the microkernel sets up for it are virtual DMA addresses that map through some IOMMU to physical DMA addresses, or whether the DMA addresses are physical DMA addresses. The addresses are simply DMA addresses, and such mappings are transparent to *DMAMemory*, and entirely maintained by the microkernel.

[0033] Device drivers can thus obtain a base address, using methods of *DMAMemory* described in detail below, which can be passed to the DMA controller, thereby allowing device drivers to provide full DMA functionality in a completely platform-independent manner.

AccessibleMemory Abstract Class

[0034] *AccessibleMemory* is what is given to drivers by bus managers, so that they can always access memory in a platform-independent manner, thus allowing them to be portable across all platforms. *AccessibleMemory* thus defines platform-independent abstract methods, such as *setByte()*. Drivers should only use the methods abstracted in *AccessibleMemory*, or methods abstracted or implemented by the super-classes of *AccessibleMemory*, which are *MainMemory*, and *Memory*. Classes below *AccessibleMemory*, and also *PhysicalMemory* and *DMAMemory*, should only be accessed by bus managers, which understand platform specifics. That way, platform specifics will not creep into driver code. *AccessibleMemory* is either accessed by a CPU issuing normal load and store instructions, hence *VirtualMemory*, or it is accessed by a CPU issuing special I/O port instructions, hence *PortIOMemory*.

PhysicalMemory Class

[0035] In this memory model, *PhysicalMemory* is not accessible, because even on systems that do not have MMUs, the abstraction is maintained that the CPUs are really accessing virtual memory (not physical memory) that happens to be mapped one-to-one to physical memory. This abstraction thus provides a means of handling systems with and without MMUs in the same general way. Since the *PhysicalMemory* class is a sub-class of the abstract class, *MainMemory*, *PhysicalMemory* provides implementations of the cache management methods abstracted in *MainMemory*.

The *PhysicalMemory* class allows a bus manager to specify a range of physical addresses, perhaps for registers of a device driver. By itself, though, since *PhysicalMemory* cannot be accessed, it is not useful. But, there are methods of *VirtualMemory* and of *DMAMemory* for mapping this physical memory into virtual memory or DMA, respectively.

VirtualMemory Abstract Class

[0036] The *VirtualMemory* abstract class contains methods for setting up mappings to *PhysicalMemory* objects (mappings ultimately set up by the microkernel). As mentioned above, even on systems with no MMUs, *VirtualMemory* objects are mapped one-to-one to *PhysicalMemory* objects. The *VirtualMemory* class also has a method for allocating a specific amount of virtual memory (also ultimately done by the microkernel), without specifying any particular physical memory. There is even a method which then allows one to obtain an array of the underlying *PhysicalMemory* objects for a *VirtualMemory* object, which is especially useful when *PhysicalMemory* objects were not passed to construct the *VirtualMemory* object in the first place. It is possible for a bus manager to set up DMA to virtual memory, since, given a *VirtualMemory* object, the array of underlying *PhysicalMemory* objects can be obtained, and then passed to a method of *DMAMemory* to set up DMA. *VirtualMemory* class also supports *lock()* and *unlock()* methods that can be used before obtaining underlying physical memory parameters before setting up a DMA operation. As set forth in detail below, there is even a *lockContiguous()* method, which makes sure the underlying physical pages mapped are contiguous, which is necessary on platforms that do not have a means, such as an IOMMU, of doing scatter-gather. It is, of course, up to the specific bus manager, which understands platform specifics, to decide what methods are needed for the specific platform.

[0037] The *VirtualMemory* abstract class also implements cache management methods (ultimately performed by the microkernel), which are abstracted by *MainMemory*. These methods, too, depending upon the coherency characteristics of the platform (also platform specifics known by the bus manager), are useful when setting up DMA.

PortIOMemory Abstract Class

[0038] The other type of *AccessibleMemory* besides *VirtualMemory* is *PortIOMemory*, which is intended for CPUs issuing special port I/O operations, such as the Intel x86 series. The *PortIOMemory* abstract class itself only implements the cache management operations abstracted by *MainMemory*. In this case, though, there is no caching, and these methods do not do much.

SwappedPortIOMemory and UnSwappedPortIOMemory Classes

[0039] The bulk of the support for *PortIOMemory* is in its two sub-classes, *SwappedPortIOMemory* and *UnSwappedPortIOMemory*, which each implement all the memory access methods abstracted by *AccessibleMemory*. As the names would imply, one does byte-swapping of the data written (to support different endi-

anness between the CPU and memory), and the other does not. It is up to the bus manager (which knows such endianness) to decide which of these two to instantiate.

[0040] The reason for having two classes, one for swapped and one for unswapped, is so it is not necessary to have a conditional that checks for whether or not memory is byte-swapped in each of the native methods. Also, the driver should not be required to know whether or not memory is swapped, since that would run counter to the need for full driver platform independence.

VirtualIOMemory and VirtualRegularMemory Abstract Classes

[0041] *VirtualMemory* is either mapped to I/O space, hence *VirtualIOMemory*, or it is not, hence *VirtualRegularMemory*. The reason for the two different abstract classes is that on some platforms accesses to I/O space may also require operations to flush out the pipeline. Both *VirtualIOMemory* and *VirtualRegularMemory* currently provide no functionality, but instead leave it to sub-classes to provide all functionality. Each of these two classes is only included in case there is ever anything common to its sub-classes that should be abstracted at a level higher.

SwappedVirtualIOMemory and UnSwappedVirtualIOMemory Classes

[0042] Similarly to *PortIOMemory*, *VirtualIOMemory* has two sub-classes, *SwappedVirtualIOMemory* and *UnSwappedVirtualIOMemory*, which each implement all the memory access methods abstracted by *AccessibleMemory*.

SwappedVirtualRegularMemory and UnSwappedVirtualRegularMemory Classes

[0043] Also similarly to *PortIOMemory*, *VirtualRegularMemory* has two sub-classes, *SwappedVirtualRegularMemory* and *UnSwappedVirtualRegularMemory*, which each implement all the memory access methods abstracted by *AccessibleMemory*.

Addresses

[0044] DMA memory, physical memory, and virtual memory are all addressed by addresses from the same *Address* abstract class. Since there are significant example of systems that have addresses with non-power-of-two number of bits, the *Address* class supports addresses with any number of bits between 1 and 64.

Unsigned Values Representing Addresses

[0045] Internally to the class, an *Address* object is

represented as an unsigned value, along with an indication of the number of bits. Since the Java programming language does not have unsigned arithmetic as part of its foundation classes, it was necessary to develop some simple efficient algorithms for addition and comparison of addresses. The *Address* abstract class is abstract mainly because it is most efficient to use different algorithms for unsigned arithmetic, depending upon the number of bits used to represent the address value. There are two sub-classes, *Address32* for addresses up to 32 bits, and *Address64*, for addresses up to 64-bits.

Platform Dependence and Platform Independence

[0046] When a bus manager constructs any of the Memory and DMA classes, it must provide a base Address and a length (also represented as an Address). To do this, it must first construct each of these Address objects, and doing so requires knowing the number of bits for that the specific type of address on the specific platform (platform specific information bus managers do know).

[0047] Drivers, on the other hand, always access AccessibleMemory objects that have been constructed by bus managers, and drivers always do so by providing an offset from the base of the AccessibleMemory object. So that drivers can be readily independent of platform specifics, such as the number of bits in a particular address, each memory access method of AccessibleMemory specifies the offset as an int. This allows memory objects up to 2 gigabytes. Although no presently known platform makes use of larger memory objects, methods can easily be added to the class to achieve even larger memory objects.

Representation of Virtual Addresses to the Microkernel

[0048] For simplicity, once an address is passed to the microkernel, it is represented (in native methods) simply as a 64-bit unsigned integer (for the value).

VirtualAddressSpace Class

[0049] Some of the machines that utilize the invention may have MMUs, so it is desirable to support multiple virtual address spaces, which allows there to be multiple virtual machines. With some prior art systems, applications always access memory in the same virtual address space as that of the drivers. But, with multiple virtual machines supported running in multiple address spaces, sometimes an application will be in a different address space from its driver. That is, a given driver will be accessing a memory object and will have synchronization methods for such access to handle the possibility of multiple threads of the same virtual machine running in that driver simultaneously. Since the virtual machines

do not have efficient mechanisms for communication that could be used to provide synchronization for instances of the same driver running in separate virtual machines, but possibly accessing the same underlying memory object (say a register that is mapped in for both), it follows that a particular driver will probably always be in the same virtual address space, while applications that use that driver may be in different address spaces.

[0050] To address the above considerations, a *VirtualAddressSpace* class is provided as part of the Memory and DMA Classes, and it has a single static method, *getVasID()*, for obtaining the virtual address space ID of the currently running thread. This ID is obtained from the microkernel. For efficiency, it is cached by *getVasID()*, so the microkernel does not have to be called upon each time.

Use of Virtual Address Space ID

[0051] When VirtualMemory objects are constructed, the virtual address space ID is one of the parameters of the constructor. This ID is eventually provided to the microkernel when the microkernel is requested to map the virtual memory. It follows that only the microkernel assigns semantics to the value of the virtual address space ID.

Device Interface Classes

[0052] The material below describes classes 50 which constitute a first package of data types, including memory classes. These classes are platform independent and can thus be imported by platform-independent Java programming language code of the operating system, such as drivers, thereby providing fully functional platform-independent drivers. Also, since they are platform-independent, these classes do not depend upon any of the platform-dependent memory classes 72 (Fig. 2).

[0053] For each class, all public fields are documented.

Address

[0054]

public abstract class Address

[0055] All memory objects of the Memory and DMA classes use Address objects for their addresses. Since there are some actual significant examples of hardware which has addresses with non-power-of-two number of bits, addresses of any number of bits between 1 and 64 are supported. As the Address class is abstract, it is extended by

- Address32, which provides support for addresses between 1 and 32 bits, and

- Address64, which provides support for addresses between 33 and 64 bits.

The Address class defines public abstract methods which are implemented by Address32 and Address64:

```
public abstract String toString();
public abstract int intValue();
public abstract long longValue();
```

[0056] In addition to the protected constructor and public methods specified below, which are common to both sub-classes, the Address abstract class also implements a package private method for checking whether or not two Addresses are of the same type (same number of bits).

Address

[0057]

protected Address (int bits)

The Address class itself only keeps track of the number of bits. The value of the address is stored in the sub-class.

getBitWidth

[0058]

```
public int getBitwidth()
```

Returns:

The number of bits (between 1 and 64) of this address.

addAddr

[0059]

```
public Address addAddr(Address other) throws
InvalidAddressException;
```

Parameters:

other - The Address to add to this Address.

Returns:

An Address whose value is the sum of this Address and the other Address, and which has the same number of bits as both this Address and the other Address.

Throws:

InvalidAddressException is thrown if the other Address does not have the same number of bits as this Address, or if the unsigned sum of the values of both Addresses overflows the maximum unsigned value than can be represented with the number of bits common to these two Addresses.

[0060] This method is included for completeness, so that a client can form an Address which is the sum of a base Address and an offset (also represented as an Address), without the client having to implement its own algorithms for unsigned arithmetic. However, since clients always access (read or write) memory objects by specifying only an offset, leaving the Memory classes to compute the address of access (sum of the base Address and offset), the clients should not really need to use this method.

[0061] This method invokes the appropriate method of the underlying Address sub-class.

addrCompare

[0062]

```
public int addrCompare(Address other) throws
InvalidAddressException;
```

Parameters:

other - The Address to compare with this address.

Returns:

-1 if the value of this Address is greater than the value of the other Address.

0- if the value of this Address is equal to the value of the other Address.

-1- if the value of this Address is less than the value of the other Address.

Throws:

InvalidAddressException is thrown if the other Address does not have the same number of bits as this Address.

[0063] Clients access (read or write) a memory object by specifying an offset (also represented as an Address). This method allows a client to determine whether this offset is valid (less than the length of the memory object), without having to implement its own algorithms for unsigned arithmetic. However, it is not absolutely necessary that the client use this method, since the Memory class access methods already do bounds checking.

[0064] This method invokes the appropriate method of the underlying Address sub-class.

Address32

[0065]

```
public class Address32 extends Address
```

[0066] The Address32 class extends the abstract Address class, and provides support for addresses whose number of bits is between 1 and 32, inclusive. The public methods specified here consist of construc-

tion methods and implementations of abstract public methods of the Address abstract class.

[0067] It should be noted that the Address32 class also implements package private methods for unsigned 32-bit addition and unsigned 32-bit comparison, which are more efficient for internal calculations than the public methods, addAddr() and addrCompare() of the Address class. These package-private methods are more efficient because they do not have to allocate the Address object to return.

Address32

[0068]

```
public Address32(int addrvalue, int bitwidth) throws
InvalidAddressException
```

```
public Address32(int addrvalue)
```

```
public Address32(long addrvalue, int bitwidth)
throws InvalidAddressException
```

Parameters:

addrvalue - The actual value of the Address. If bitwidth is specified, and it is less than 32, only the least significant bitwidth number of bits is used for the value of the Address.
bitwidth - The number of bits (between 1 and 32) of the Address.

Throws:

Invalid Address Exception is throws if a bitwidth is specified which is not between 1 and 32, inclusive.

[0069] The first of these constructors allows the client to create a general type of address of between 1 and 32 bits. The second constructor is provided for convenience to construct the more common 32-bit Address. The third constructor is provided for convenience, when the value for the Address is a long. In this case, though, only the least significant 32-bits (or less) are used.

to String

[0070]

```
public String to String()
```

Returns:

A string representing this Address32 object.

intValue

[0071]

```
public int intvalue()
```

Returns:

The least significant 32 bits of the address value.

long Value

[0072]

```
public long Value()
```

Returns:

The least significant 64 bits of the address value.

10 Address64

[0073]

```
public class Address 64 extends Address
```

15

[0074] The Address64 class extends the abstract Address class, and provides support for addresses whose number of bits is between 33 and 64, inclusive. The public methods specified here consist of construction methods and implementations of abstract public methods of the Address abstract class.

20

[0075] It should be noted that the Address64 class also implements package private methods for unsigned 64-bit addition and unsigned 64-bit comparison, which are more efficient for internal calculations than the public methods, addAddr() and addrCompare() of the Address class. These package-private methods are more efficient because they do not have to allocate the Address object to return.

25

30

Address64.

[0076]

35

```
public Address64(long addrvalue, int bitwidth)
throws InvalidAddressException
```

```
public Address64(long addrvalue)
```

40

Parameters:

addrvalue - The actual value of the Address. If bitwidth is specified, and it is less than 64, only the least significant bitwidth number of bits is used for the value of the Address.

45

bitwidth - The number of bits (between 33 and 64) of the Address.

Throws:

InvalidAddressException is throws if a bitwidth is specified which is not between 33 and 64, inclusive.

50

[0077] The first of these constructors allows the client to create a general type of address of between 33 and 64 bits. The second constructor is provided for convenience to construct the more common 64-bit Address.

55

toString

[0078]

public String toString

Returns:

A string representing this Address64 object.

intValue

[0079]

public int intValue()

Returns:

The least significant 32 bits of the address value.

longValue

[0080]

public long longvalue()

Returns:

The least significant 64 bits of the address value.

AddressSpace

[0081]

public abstract class AddressSpace

[0082] Every bus manager manages a set of address spaces. The memory classes themselves have three address spaces (see VirtualAddressSpace, PhysicalAddressSpace, and DMAAddressSpace) which are all logically part of the particular bus manager known as the PlatformBusManager, the highest level bus manager. Besides these address spaces, there are other address spaces that other bus managers can instantiate. All such address spaces are sub-classes of the AddressSpace abstract class.

AddressSpace

[0083]

protected AddressSpace (ExpansionBus b, String n)

Parameters:

b - Specifies the bus that has this particular address space.

n-A string naming this address space.

[0084] This constructor simply initializes protected variables with the supplied values.

[0085] A protected initialization routine is also pro-

vided so that the sub-class can initialize the bit-size of the address space and an id for the address space. Both of these initializers might be obtained by the sub-class via native methods. In the case where there may be multiple address spaces of a particular type (e.g. virtual), the id obtained from the microkernel uniquely specifies one, so that it can be specifically identified. Using the initializer for the number of bits of the address space, the AddressSpace class computes and stores Address objects representing the lowest and highest possible addresses.

getId

[0086]

public int getId()

Returns:

The id that specifies which particular address space (see above).

getName

[0087]

public String getName()

Returns:

The String used to initialize this address space in the constructor.

getExpansionBus

[0088]

public String getName()

Returns:

The bus used to initialize this address space in the constructor.

getLowestAddress

[0089]

public String getLowestAddress()

Returns:

The lowest address value possible in this address space.

getHighestAddress

[0090]

public String getHighestAddress()

Returns:

The highest address value possible in this address

space.

MemoryConstraints

[0091]

public class MemoryConstraints

[0092] When a bus manager creates a memory object, it may need to specify constraints of both the allocation of the memory object, and subsequent accesses to the memory object by a driver. For example, some DMA controllers may have a limited range of addresses they can emit, thus necessitating that a DMA object (a memory object) be allocated only in that address range. As an example of access constraints, some registers may only be accessible as bytes, such that any other access may cause some sort of system failure.

[0093] The MemoryConstraints object encapsulates such allocation and access constraints, so that they can be passed together to a constructor of a memory object.

[0094] The MemoryConstraints class contains a constructor used to specify all the various constraints, and several methods for obtaining and setting each of the various constraints. This class also contains package private methods for checking whether a requested access satisfies the constraints, and whether a specified allocation matches the constraints.

[0095] The following constants are defined:

```
public static final int ALIGN_BYTE = 0;
public static final int ALIGN_SHORT = 1;
public static final int ALIGN_INT = 2;
public static final int ALIGN_LONG = 3;
public static final int ALIGN_CACHE = 4;
public static final int ALIGN_PAGE = 5;
public static final int READONLY = 0;
public static final int READWRITE = 1;
```

MemoryConstraints

[0096]

```
public MemoryConstraints (Address alloc_min_a,
Address alloc_max_a, int alloc_align, int cache,
boolean lock, boolean waitformem, boolean swapped,
boolean io, boolean virtual, int access_min_s,
int access_max_s, int access_align, int rw) throws InvalidAddressException
```

Parameters:

alloc_min_a - The minimum address for allocation.
alloc_max_a - The maximum address for allocation.

alloc_align - The minimum number of least significant bits of the base address of an allocation that must be all 0s. A value of 0 thus indicates there are

no alignment constraints for allocations. The maximum value supported of alloc_align is 16 (64K byte alignment).

cache - Indicates caching of the memory (either MainMemory.CACHE_MODE_DEFAULT, MainMemory.CACHE_MODE_INHIBITED, MainMemory.CACHE_MODE_WRITE_THROUGH, MainMemory.CACHE_MODE_COPY_BACK).

lock - Indicates whether or not the mappings for the memory are to be locked.

swapped - If true, indicates a type of memory is to be constructed, where byte-swapping is to be done.

io - If true, indicates a type of memory is to be constructed that is I/O memory (either PortIOMemory or VirtualIOMemory).

virtual - If true, indicates a type of memory is to be constructed that is VirtualMemory

waitformem - Indicates whether or not to wait indefinitely for the allocation of the memory, if there is a resource problem.

access_min_s - The log (base2) of the minimum size in bytes of memory that can be accessed. Thus, 0 indicates that a single byte can be accessed.

access_max_s - The log (base 2) of the maximum size in bytes of memory that can be accessed (up to 3).

access_align - The minimum number of least significant bits of an access address that must be all 0s. A value of 0 thus indicates that there are no alignment constraints for accesses. The maximum value supported of access_align is 16 (64K byte alignment).

rw - READONLY (0) indicates read-only access, while READWRITE(1) indicates read-write access.

Throws:

InvalidAddressException is thrown if alloc_min_a and alloc_max_a do not agree on their number of bits, or if the value of alloc_min_a is not less than the value of alloc_max_a

[0097] The maximum size allocation alignment supported is ALIGN_PAGE (5), and the minimum is, of course, ALIGN_BYTE(0), so if the specified alloc_align is greater than ALIGN_PAGE (5), it is set to 5, or if the specified alloc_align is less than ALIGN_BYTE(0), it is set to 0.

[0098] If the specified access_min_s is less than 0, it is set to 0 (1 byte), or if the specified access_minds is greater than 3, it is set to 3 (8 bytes, a long). Similarly, if the specified access_max_s is less than 0, it is set to 0 (1 byte), or if the specified access_max_s is greater than 3, it is set to 3 (8 bytes, a long). If access_max_a is still less than access_min_a, access_max_a is set to the value of access_min_a.

[0099] The maximum size access alignment supported is ALIGN_LONG(3), and the minimum sup-

ported is, of course, ALIGN_BYTE(0); so if the specified access_align is greater than ALIGN_LONG(3), it is set to 3, or if the specified access_align is less than ALIGN_BYTE(0), it is set to 0.

[0100] If rw is not set to READONLY, it is set to READ-
WRITE.

getAllocMinAddr

[0101]

```
public Address getAllocMinAddr()
```

Returns:

The minimum allocation Address, alloc_min_a, which was specified by either the constructor, or by the setAllocAdrs() method.

getAllocMaxAddr

[0102]

```
public Address getAllocMaxAddr()
```

Returns:

The maximum allocation Address, alloc_max_a, which was specified by either the constructor, or by the setAllocAdrs() method.

setAllocAdrs

[0103]

```
public void setAllocAdrs(Address alloc_min_a,  
Address alloc_max_a) throws InvalidAddressEx-  
ception
```

Parameters:

alloc_min_a - The minimum address for allocation.
alloc_max_a - The maximum address for allocation.

Throws:

InvalidAddressException is thrown if alloc_min_a and alloc_max_a do not agree on their number of bits, or if the value of alloc_min_a is not less than the value of alloc_max_a.

getAllocAlign

[0104]

```
public int getAllocAlign()
```

Returns:

The alignment for allocations, alloc_align, as specified by either the constructor, or by the setAllocAlign() method. Note that this value may be different from

the value specified (see descriptions of the constructor or the setAllocAlign method).

setAllocAlign

[0105]

```
public void setAllocAlign(int alloc_align)
```

Parameters:

alloc_align - The minimum number of least significant bits of the base address of an allocation that must be all 0s. A value of 0 thus indicates there are no alignment constraints for allocations. The maximum value supported of alloc_align is 16 (64K byte alignment).

[0106] The maximum size allocation alignment supported is ALIGN_PAGE (5), and the minimum is, of course, ALIGN_BYTE(0), so if the specified alloc_align is greater than ALIGN_PAGE (5), it is set to 5, or if the specified alloc_align is less than ALIGN_BYTE(0), it is set to 0.

getCacheMode

[0107]

```
public int getCacheMode()
```

Returns:

The cache mode, cache, as specified by either the constructor, or by the setMiscAlloc() method.

getLocked

[0108]

```
public boolean getLocked()
```

Returns:

The indication of whether mappings are to be locked upon allocation, lock, as specified by either the constructor, or by the setMiscAlloc() method.

getWaitforMem

[0109]

```
public boolean getwaitFormMem()
```

Returns:

The indication of whether or not to wait for memory indefinitely during allocation, waitformem, as specified by either the constructor, or by the setMiscAlloc() method.

getSwapped**[0110]**

```
public int getSwapped()
```

Returns:

The argument, swapped, as specified by either the constructor, or by the setMiscAlloc() method.

getIO**[0111]**

```
public int getIO()
```

Returns:

The argument, io, as specified by either the constructor, or by the setMiscAlloc() method.

getVirtual**[0112]**

```
public int getVirtual()
```

Returns:

The argument, virtual, as specified by either the constructor, or by the setMiscAlloc() method.

setMiscAlloc**[0113]**

```
public void setMiscAlloc(int cache, boolean lock,
boolean waitformem, boolean swapped, boolean io,
boolean virtual)
```

Parameters:

cache - Indicates caching of the memory (either MainMemory.CACHE_MODE_DEFAULT, MainMemory.CACHE_MODE_INHIBITED, MainMemory.CACHE_MODE_WRITE_THROUGH, or MainMemory.CACHE_MODE_COPY_BACK).

lock - Indicates whether or not the mappings for the memory are to be locked.

waitformem - Indicates whether or not to wait indefinitely for the allocation of the memory, if there is a resource problem.

swapped - If true, indicates a type of memory is to be constructed, where byte-swapping is to be done.

io - If true, indicates a type of memory is to be constructed that is I/O memory (either PortIOMemory or VirtualIOMemory).

virtual - if true, indicates a type of memory is to be constructed that is VirtualMemory.

getAccessMinSize**[0114]**

```
public int getAccessMinSize()
```

Returns:

The minimum size for accesses, access_min_s, as specified by either the constructor, or by the setAccessSizes() method. Note that this value may be different from the value specified (see descriptions of the constructor or the setAccessSizes() method).

getAccessMaxSize

```
public int getAccessMaxSize()
```

[0115]

Returns:

The maximum size for accesses, access_max_s, as specified by either the constructor, or by the setAccessSizes() method. Note that this value may be different from the value specified (see descriptions of the constructor or the setAccessSizes() method).

setAccessSizes**[0116]**

```
public void setAccessSizes(int access_min_s, int
access_max_s)
```

Parameters:

access_min_s - The log (base 2) of the minimum size in bytes of memory that can be accessed. Thus, 0 indicates that a single byte can be accessed.

access_max_s - The log (base 2) of the maximum size in bytes of memory that can be accessed (up to 3).

[0117] If the specified access_min_s is less than 0, it is set to 0 (1 byte), or if the specified access_min_s is greater than 3, it is set to 3 (8 bytes, a long). Similarly, if the specified access_max_s is less than 0, it is set to 0 (1 byte), or if the specified access_max_s is greater than 3, it is set to 3 (8 bytes, a long). If access_max_a is still less than access_min_a, access_max_a is set to the value of access_min_a.

getAccessAlign**[0118]**

```
public void getAccessAlign()
```

Returns:

The alignment for accesses, `access_align`, as specified by either the constructor, or by the `setAccessAlign()` method. Note that this value may be different from the value specified (see descriptions of the constructor or the `setAccessAlign` method).

setAccessAlign

[0119]

```
public void setAccessAlign(int access_align)
```

Parameters:

`access_align` - The minimum number of least significant bits of an access address that must be all 0s. A value of 0 thus indicates that there are no alignment constraints for accesses. The maximum value supported of `access_align` is 16 (64K byte alignment).

[0120] The maximum size access alignment supported is `ALIGN_LONG(3)`, and the minimum supported is, of course, `ALIGN_BYTE(0)`, so if the specified `access_align` is greater than `ALIGN_LONG(3)`, it is set to 3, or if the specified `access_align` is less than `ALIGN_BYTE(0)`, it is set to 0.

getReadWrite

[0121]

```
public int getReadwrite()
```

Returns:

The mode of access, either `READONLY`, or `READWRITE`. Note that this value may be different from the value specified (see descriptions of the constructor or the `setReadWrite` method).

setReadWrite

[0122]

```
public void setReadwrite(int rw)
```

Parameters:

`rw` - `READONLY` indicates read-only access, while `READWRITE` indicates readwrite access.

[0123] If `rw` is not set to `READONLY`, it is set to `READWRITE`.

Memory

[0124]

```
public abstract class Memory
```

[0125] All memory objects of the `Memory` and `DMA` classes are ultimately sub-classes of the abstract class `Memory`. The `Memory` abstract class defines two abstract methods:

```
public abstract String toString();
```

[0126] In addition to the fundamental public methods specified below that are common to all memory objects, the `Memory` class also implements package private methods for checking validity of an access (specified by offset within a memory object, size, mode (read or write), and count in bytes), and for obtaining the address of an access, given an offset of the memory object.

Memory

[0127]

protected `Memory` (`Address ba`, `Address len`, `MemoryConstraints mc`) throws `InvalidAddressException`

protected `Memory` ()

protected `Memory` (`Memory superrange`, `Address offset`, `Address len`) throws `InvalidAddressException`, `AllocationException`

Parameters:

`ba` - The base `Address` of the `Memory` object.
`len` - The length in bytes of the `Memory` object.
`mc` - Constraints of the memory object

Throws:

`InvalidAddressException` is thrown if an overflow occurs when adding the base `Address`, `ba`, to the length, `len`.

`AllocationException` is thrown, in the case of the third protected constructor, if the super-range `Memory` object is invalid.

[0128] This first protected constructor simply calculates the end `Address` (`ba + len`), saves all the parameters, along with the calculated end `Address` in protected variables of this class, initializes the memory to valid, and allocates protected internal data structures.

[0129] The second protected constructor is used in the case where the underlying `Memory` object cannot be initialized when first constructed (e.g. `VirtualMemory` whose base address will not be known until the memory has been mapped in). A protected method is provided for later initializing the `Memory` object, thus completing what would have been initialized had the first protected constructor been used.

[0130] The third protected constructor is used in the case where the memory object being constructed is a

sub-range of an existing memory object. The memory base Address and end Address are calculated based upon the offset to the superrange Memory object, and the new Memory object is associated as a sub-range of the superrange and initialized to valid. Protected internal data structures area also allocated.

getMemBaseAddr

[0131]

```
public Address getMemBaseAddr()
```

Returns:

The base Address of the Memory object if the memory is valid, and null otherwise.

getMemLength

[0132]

```
public Address getMemLength()
```

Returns:

The length of the Memory object in bytes if the memory is valid, and null otherwise.

getMemEndAddr

[0133]

```
public Address getMemEndAddr ()
```

Returns:

The end Address of the Memory object if the memory is valid, and null otherwise.

getMemValid

[0134]

```
public boolean getMemvalid()
```

Returns:

The indication of whether or not the memory is valid.

getMemConstraints

[0135]

```
public MemoryConstraints getMemConstraints ()
```

Returns:

The constraints of the memory if the memory is valid, otherwise null.

getSuperRange

[0136]

```
public Memory getSuperRange()
```

Returns:

The Memory object which is a super-range of this Memory object, if it is a sub-range; otherwise returns null.

getSubRanges

[0137]

```
public Vector getSubRanges()
```

Returns:

The a Vector of Memory objs which are sub-ranges of this Memory object. If there are no sub-ranges or if this Memory object is invalid, null is returned.

MemoryDescriptor

[0138]

```
public class MemoryDescriptor extends Memory
```

[0139] All real instantiated memory objects are ultimately platform-dependent and are thus constructed by code in the sun.javaos.memory package. Sometimes, though, it is useful for drivers or bus managers to be able to instantiate a memory object, which is a descriptor for memory that will be allocated later. The MemoryDescriptor class, as an instantiable sub-class of Memory, meets that need.

[0140] Associated with a MemoryDescriptor object is an address space, managed by some bus manager, within which this descriptor of memory is defined.

MemoryDescriptor

[0141]

```
public MemoryDescriptor(Address ba, Address len,
MemoryConstraints mc, AddressSpace as, String
name) throws InvalidAddressException
```

```
public MemoryDescriptor(Address ba, Address len,
AddressSpace as, String name) throws InvalidAd-
dressException
```

Parameters:

ba - The base Address of the Memory object.

len - The length in bytes of the Memory object.

mc - Constraints of the memory object.

as - The AddressSpace in which this memory object lives.

name - Some name that is used to identify this memory object. This is useful when a bus manager is allocating an array of MemoryDescriptors, specifying all the various memory objects that a driver needs to have mapped, since it allows the driver to identify each of the memory objects.

Throws:

InvalidAddressException is thrown if the constructor of the super-class (Memory) throws InvalidAddressException. In other words, this constructor does not catch that exception.

[0142] The first constructor simply calls the constructor of the super-class, Memory, using all but the AddressSpace argument, and then stores the AddressSpace in a private variable.

[0143] The second constructor is provided for convenience, and is the same as the first constructor, except that null is passed to the Memory constructor for the mc (MemoryConstraints) parameter.

toString

[0144]

public String toString()

Returns:

A string representing this MemoryDescriptor object.

getAddressSpace

[0145]

public AddressSpace getAddressSpace ()

Returns:

The AddressSpace in which this descriptor of Memory exists, as specified in the constructor.

getName

[0146]

public String getName()

Returns:

The string naming this particular MemoryDescriptor, as specified in the constructor.

MainMemory

[0147]

public abstract class MainMemory extends Memory

[0148] All memory objects which represent memory

that can be accessed ultimately by a CPU (everything except DMAMemory, the already described super-class Memory, and MainMemory itself) is a sub-class of the abstract class, MainMemory.

[0149] MainMemory defines the following abstract methods:

public abstract void setCacheMode(int mode);

public abstract int getCacheMode();

public abstract void flushCache();

[0150] MainMemory also defines the following constants to be used as parameters for implementations of setCacheMode() and return values of implementations of getCacheMode():

public static final int
CACHE_MODE_DEFAULT=0x0;

public static final int
CACHE_MODE_AMBIGUOUS=0x1;

public static final int
CACHE_MODE_INHIBITED=0x2;

public static final int
CACHE_MODE_WRITE_THROUGH=0x3;

public static final int
CACHE_MODE_COPY_BACK=0x4;

[0151] The only methods implemented for this class are protected constructors. No public methods are implemented.

MainMemory

[0152]

protected MainMemory (Address ba, Address len, MemoryConstraints mc) throws InvalidAddressException

protected MainMemory ()

protected MainMemory (MainMemory superrange, Address offset, Address len) throws InvalidAddressException, AllocationException

Parameters:

ba - The base Address of the Memory object.
len - The length in bytes of the Memory object.
mc - Constraints of the memory object.

Throws:

InvalidAddressException is thrown if an overflow

occurs when adding the base Address, ba, to the length, len.

AllocationException is thrown, in the case of the third protected constructor, if the super-range Memory object is invalid.

[0153] Each of these protected constructors simply calls the corresponding protected constructor of the super-class, Memory.

AccessibleMemory

[0154]

```
public abstract class AccessibleMemory extends
MainMemory
```

[0155] All memory objects of the Memory and DMA classes which represent memory that can be accessed directly by a CPU (every sub-class of MainMemory except PhysicalMemory and AccessibleMemory itself) is a sub-class of the abstract class, AccessibleMemory.

[0156] Platform-independent code, such as drivers, always access Memory via the methods of AccessibleMemory, not knowing or caring what the underlying memory object may be.

[0157] AccessibleMemory defines the following protected abstract methods, which are internal implementations of the publicly defined access methods:

```
protected abstract int ncksum(long adr, long len);
```

```
protected abstract void nSetByte(long adr, byte x);
```

```
protected abstract void nSetShort(long adr, short
x);
```

```
protected abstract void rSetInt(long adr, int x);
```

```
protected abstract void nsetLong(long adr, long x);
```

```
protected abstract void nSetBytes(long adr, byte
bytes[],int bytesOffset, int length);
```

```
protected abstract void nsetIntArray(long adr, byte
bytes[],int bytesOffset, int length_int);
```

```
protected abstract byte nGetByte(long adr);
```

```
protected abstract short nGetShort(long adr);
```

```
protected abstract int nGetInt(long adr);
```

```
protected abstract long nGetLong(long adr);
```

```
protected abstract void nGetBytes(long adr, byte
bytes[],int bytesOffset, int length);
```

```
protected abstract void nGetIntArray(long adr, byte
bytes[],int bytesOffset, int length_int);
```

```
protected abstract int nRegisterSource(long adr1,
long len);
```

```
protected abstract void nCopy(long adr2, long len,
int copyID);
```

10 [0158] There are no other non-public methods defined, besides the constructor.

[0159] AccessibleMemory does not implement any of the abstract methods of the two classes above it in the hierarchy, MainMemory and Memory.

15 [0160] It should be noted that each of the access methods below checks for invalid accesses, and if there is a violation, a system message is logged, and the access is not performed. The reason for not simply having these methods throw an exception is that then all current calls in various drivers - and there are very many - would need to be modified to declare the exception thrown, when drivers are ported to the new Memory classes.

AccessibleMemory

[0161]

```
protected AccessibleMemory(Address ba, Address
len, MemoryConstraints mc) throws InvalidAd-
dressException
```

```
protected AccessibleMemory ()
```

```
protected AccessibleMemory (AccessibleMemory
superrange, Address offset, Address len) throws
InvalidAddressException, AllocationException
```

Parameters:

ba - The base Address of the Memory object.
len - The length in bytes of the Memory object.
mc - Constraints of the Memory object.

Throws:

InvalidAddressException is thrown if an overflow occurs when adding the base Address, ba, to the length, len.

AllocationException is thrown, in the case of the third protected constructor, if the super-range Memory object is invalid.

[0162] Each of these protected constructors simply calls the corresponding protected constructor of the super-class, MainMemory.

55

setAddress**[0163]**

```
public void setAddress(int offset, Address value)
```

5

Parameters:

offset - Offset from the base of the memory object to the place where the access (in this case a write) is to be done.

value - The Address whose value is to be written at the access location.

Note:

The access is not done and a system message is logged if the memory is not valid (defined in Memory class), if the offset (plus size of access) is beyond the end of the memory object, or if any of the constraints defined in the constraints of the super-class Memory object is violated.

15

20

cksum**[0164]**

```
public int cksum(int offset, int len)
```

25

Parameters:

offset - Offset from the base of the memory object to the place where the checksum computation is to start.

len - The length in bytes over which the checksum is to be computed.

30

Returns:

The computed checksum (modulo 0x10000), or 0 if the access is not allowed (see Note).

35

Note:

The checksum is not done (but 0 is returned) and a system message is logged if the memory is not valid (defined in Memory class), or if the offset (plus length, len, of access) is beyond the end of the memory object.

40

45

setByte**[0165]**

```
public void setByte(int offset, byte x)
```

50

Parameters:

offset - Offset from the base of the memory object to the place where the access (in this case a write) is to be done.

x - The value to be written at the access location.

55

Note:

The access is not done and a system message is logged if the memory is not valid (defined in Memory class), if the offset (plus size of access) is beyond the end of the memory object, or if any of the constraints defined in the constraints of the super-class Memory object is violated.

setShort**[0166]**

```
public void setShort(int offset, short x)
```

Parameters:

offset - Offset from the base of the memory object to the place where the access (in this case a write) is to be done.

x - The value to be written at the access location.

Note:

The access is not done and a system message is logged if the memory is not valid (defined in Memory class), if the offset (plus size of access) is beyond the end of the memory object, or if any of the constraints defined in the constraints of the super-class Memory object is violated.

setInt**[0167]**

```
public void setInt(int offset, int x)
```

Parameters:

offset - Offset from the base of the memory object to the place where the access (in this case a write) is to be done.

x - The value to be written at the access location.

Note:

The access is not done and a system message is logged if the memory is not valid (defined in Memory class), if the offset (plus size of access) is beyond the end of the memory object, or if any of the constraints defined in the constraints of the super-class Memory object is violated.

setLong**[0168]**

```
public void setLong(int offset, long x)
```

Parameters:

offset - Offset from the base of the memory object to the place where the access (in this case a write) is to be done.

x - The value to be written at the access location.

Note:

The access is not done and a system message is logged if the memory is not valid (defined in Memory class), if the offset (plus size of access) is beyond the end of the memory object, or if any of the constraints defined in the constraints of the super-class Memory object is violated.

setBytes

[0169]

```
public void setBytes(int offset, byte bytes[], int
bytesOffset, int length)
```

Parameters:

offset - Offset from the base of the memory object to the place where the access (in this case a write) is to be done.

bytes - The array of bytes to be written to the memory.

bytesOffset - The index within the array, bytes, to begin.

length - The number of bytes to write.

Note:

The access is not done and a system message is logged if the memory is not valid (defined in Memory class), if bytesOffset is negative, if length is not greater than 0, if the offset (plus total size of access) is beyond the end of the memory object, or if any of the constraints defined in the constraints of the super-class Memory object is violated.

setIntArray

[0170]

```
public void setIntArray(int offset, byte bytes[], int
bytesOffset, int length_int)
```

Parameters:

offset - Offset from the base of the memory object to the place where the access (in this case a write) is to be done.

bytes - The array of bytes to be written (as integers) to the memory.

bytesOffset - The index within the array bytes to begin.

length_int - The number of ints to write.

Note:

The access is not done and a system message is logged if the memory is not valid (defined in Memory class), if bytesOffset is negative, if length_int is not greater than 0, if the offset (plus total size of

access) is beyond the end of the memory object, or if any of the constraints defined in the constraints of the super-class Memory object is violated.

5 **getByte**

[0171]

```
public byte getByte(int offset)
```

Parameters:

offset - Offset from the base of the memory object to the place where the access (in this case a read) is to be done.

Returns:

The value read from the access location.

Note:

The access is not done and a system message is logged if the memory is not valid (defined in Memory class), if the offset (plus size of access) is beyond the end of the memory object, or if any of the constraints defined in the constraints of the super-class Memory object is violated.

getShort

[0172]

```
public short getShort(int offset)
```

Parameters:

offset - Offset from the base of the memory object to the place where the access (in this case a read) is to be done.

Returns:

The value read from the access location.

Note:

The access is not done and a system message is logged if the memory is not valid (defined in Memory class), if the offset (plus size of access) is beyond the end of the memory object, or if any of the constraints defined in the constraints of the super-class Memory object is violated.

getInt

[0173]

```
public int getInt(int offset)
```

Parameters:

offset - Offset from the base of the memory object to the place where the access (in this case a read) is to be done.

Returns:

The value read from the access location.

Note:

The access is not done and a system message is logged if the memory is not valid (defined in Memory class), if the offset (plus size of access) is beyond the end of the memory object, or if any of the constraints defined in the constraints of the super-class Memory object is violated.

getLong

[0174]

```
public long getLong(int offset)
```

Parameters:

offset - Offset from the base of the memory object to the place where the access (in this case a read) is to be done.

Returns:

The value read from the access location.

Note:

The access is not done and a system message is logged if the memory is not valid (defined in Memory class), if the offset (plus size of access) is beyond the end of the memory object, or if any of the constraints defined in the constraints of the super-class Memory object is violated.

getBytes

[0175]

```
public void getBytes(int offset, byte bytes[], int
bytesOffset, int length)
```

Parameters:

offset - Offset from the base of the memory object to the place where the access (in this case a read) is to be done.

bytes - The array of bytes to be read from memory.

bytesOffset - The index within the array bytes to begin.

length - The number of bytes to read.

Note:

The access is not done and a system message is logged if the memory is not valid (defined in Memory class), if bytesOffset is negative, if length is not greater than 0, if the offset (plus total size of access) is beyond the end of the memory object, or if any of the constraints defined in the constraints of the super-class Memory object is violated.

getIntArray

[0176]

```
public void getIntArray(int offset, int ints[], int
intsOffset, int length_int)
```

Parameters:

offset - Offset from the base of the memory object to the place where the access (in this case a read) is to be done.

ints-The array of ints to be written to the memory.

intsOffset - The index within the array ints to begin.

length_int-The number of ints to write.

Note:

The access is not done and a system message is logged if the memory is not valid (defined in Memory class), if intsOffset is negative, if length_int is not greater than 0, if the offset (plus total size of access) is beyond the end of the memory object, or if any of the constraints defined in the constraints of the super-class Memory object is violated.

copy

[0177]

```
public void copy(int offset, AccessibleMemory dest,
int destOffset, int length)
```

Parameters:

offset - Offset from the base of the memory object to the place where the access (in this case a read) is to be done.

dest - The memory object which contains the destination of the copy.

destOffset - Offset (expressed as either an Address, a short, an int, or a long) from the base of the memory object to the place where the access (in this case a write) is to be done.

Note:

The access is not done and a system message is logged if the memory (either source or destination) is not valid (defined in Memory class), if the offset (plus total size of access for both source and destination) is beyond the end of the memory object, if any of the constraints defined in the constraints of the super-class Memory object is violated, or if the two Memory objects are not the same basic type.

Platform Interface Classes

[0178] The material below describes classes 72 which constitute a second package of memory classes. Unlike memory classes 50 which are platform-independent, the classes of memory package 72 are platform-dependent. These classes are primarily used by platform-dependent client code of, for example, bus manag-

ers. These classes depend upon and extend the platform-independent classes of `javax.system.memory`.

PhysicalAddressSpace

[0179]

```
public final class PhysicalAddressSpace extends
    AddressSpace
```

[0180] The `PhysicalAddressSpace` is not really used by the memory classes themselves, since it is assumed that a machine only has one physical address space, but it is provided for completeness, and for use outside of the memory classes, with a `MemoryDescriptor`, by various bus managers.

[0181] Besides the constructor, the following public string is defined:

```
public final static String name = "physical";
```

PhysicalAddressSpace

[0182]

```
public PhysicalAddressSpace (ExpansionBus b)
```

[0183] This constructor first calls the constructor of the super-class, `AddressSpace`, specifying the `ExpansionBus` argument and the `String` defined above, so that they may be used by the public methods, `getExpansionBus()` and `getName()` of `AddressSpace`. Then it calls the `getAddressSize()` and the `alloc_id()` native methods (see below) and supplies them to an protected initialization routine of `AddressSpace`, so that they may then be used by the public methods, `getLowestAddress()`, `getHighestAddress()`, and `getID()` of `AddressSpace`.

[0184] The `PhysicalAddressSpace` defines the following native methods, implemented in the microkernel, which provide the size of addresses (in bits) and an ID for the address space, respectively.

```
private native int getAddressSize();
```

```
private native int alloc_id();
```

PhysicalMemory

[0185]

```
Public class PhysicalMemory extends MainMemory
```

[0186] The `PhysicalMemory` class allows bus managers to obtain access to objects at specific physical addresses (such as device registers which are listed to be at a specific physical address in the Java system Database (JSD)) and then map them into virtual memory (see `VirtualMemory` constructor methods) so they

can be accessed by device drivers (via the `AccessibleMemory` access methods). It is also possible for bus managers to set up DMA to a specified array of `PhysicalMemory` objects (see `DMAMemory` constructor methods). In fact, a bus manager can even set up DMA to a range of memory mapped into virtual memory, by first locking the virtual to physical mappings (see `lock()` method of `VirtualMemory`), then calling the appropriate method of `VirtualMemory` (see `getPhysicalMemory()` method) to obtain an array of `PhysicalMemory` objects, and then passing this array to the appropriate `DMAMemory` constructor. The reverse is also possible, where the `DMAMemory` object is first constructed, and then the underlying `PhysicalMemory` objects obtained in order to construct a `VirtualMemory` object (accessible via `AccessibleMemory`) which maps to the same `PhysicalMemory` objects.

[0187] Although `PhysicalMemory` is very useful in conjunction with `VirtualMemory` (which is `AccessibleMemory`) or in conjunction with `DMAMemory`, it is not usable by itself. This is true even on a system which has no MMU, where a CPU accesses physical memory directly. Even in that case, the `PhysicalMemory` object is not an object that is accessed, but when that `PhysicalMemory` object is mapped to virtual memory, the `VirtualMemory` object can be accessed, and its virtual addresses happen to be the same as the underlying physical addresses. In this way, semantically, drivers or any other Java clients always access `VirtualMemory` objects, and never access `PhysicalMemory` objects.

[0188] The following constructors and public methods are defined:

PhysicalMemory

[0189]

```
public PhysicalMemory (Address basePA, Address
    len) throws InvalidAddressException
```

Parameters:

basePA - The base physical address of the physical memory to be allocated.

len - The length in bytes of physical memory to be allocated.

Throws:

`InvalidAddressException` is thrown if the constructor called for the super-class (`MainMemory`) throws `InvalidAddressException`.

```
private PhysicalMemory(PhysicalMemory super-
    range, Address offset, Address len) throws Invali-
    dAddressException, AllocationException
    superRange - the super-range PhysicalMemory
    object.
```

offset - The offset within the current `PhysicalMemory` object for the new sub-range.

len - The length of the new sub-range.

Throws:

InvalidAddressException is thrown if either the offset or the offset + newLength is beyond the range of the PhysicalMemory object from which the sub-range is being obtained.

AllocationException is thrown if the superrange Memory object is no longer valid.

[0190] The first constructor is public and can be used by a bus manager to allocate a PhysicalMemory object. This allocation does not actually involve anything done in the microkernel corresponding to the actual physical memory. Rather, the PhysicalMemory object is merely a descriptor for physical memory, and used for construction of a VirtualMemory (actually an instantiable subclass of VirtualMemory) or DMAMemory object. It is at the time that such a VirtualMemory or DMAMemory object is constructed that the microkernel verifies the PhysicalMemory object.

[0191] This first constructor does not take a MemoryConstraints object as a parameter. Basically the reason is as follows. There are two types of constraints specified in a MemoryConstraints object, allocation constraints and access constraints. Allocation constraints do not apply, since nothing is actually allocated, and since the caller specifies the base address. Access constraints do not apply either, since PhysicalMemory objects are not accessible.

[0192] The second constructor is private, and is used internally by the getSubRange() method. It merely invokes the super-class (MainMemory) constructor for sub-ranges.

toString

[0193]

```
public String toString()
```

Returns:

A string representing this PhysicalMemory object.

getSubRange

[0194]

```
public PhysicalMemory getSubRange(Address offset, Address newLength) throws InvalidAddressException
```

Parameters:

offset - The offset within the current PhysicalMemory object for the new sub-range.

newLength - The length of the new sub-range.

Returns:

A PhysicalMemory object which is the specified sub-range.

Throws:

InvalidAddressException is thrown if either the offset or the offset + newLength is beyond the range of the PhysicalMemory object from which the sub-range is being obtained.

AllocationException is thrown if the superrange Memory object is no longer valid.

[0195] This method merely invokes the private constructor of this PhysicalMemory object.

getCacheMode

[0196]

```
public int getCacheMode()
```

Returns:

The cache mode, either MainMemory.CACHE_MODE_INHIBITED, MainMemory.CACHE_MODE_WRITE_THROUGH, or MainMemory.CACHE_MODE_COPY_BACK.

[0197] Note that this method allows a bus manager to determine how the physical memory is cached if the PhysicalMemory object has just been constructed, or if CACHE_MODE_DEFAULT has been specified by setCacheMode().

setCacheMode

[0198]

```
public void setCacheMode(int mode)
```

Parameters:

mode - Set to one of: MainMemory.CACHE_MODE_DEFAULT, MainMemory.CACHE_MODE_INHIBITED), MainMemory.CACHE_MODE_WRITE_THROUGH, or MainMemory.CACHE_MODE_COPY_BACK.

[0199] This method allows the bus manager, which should know about cache coherency of various components of the specific platform it supports, to control the caching of physical memory. For example, on platforms where a physical memory cache is not coherent with DMA accesses, it would be necessary for the bus manager to temporarily set the cache mode to CACHE_MODE_INHIBITED, while DMA transfers are set up.

[0200] Specifically, this method requests the following, depending upon the specific value of mode:

CACHE_MODE_DEFAULT - Cache this range of physical memory according to what the default is for this type of physical memory on this platform. Note that, even on platforms that have physical address

caches, some physical memory (e.g. I/O registers) may need to be uncached to be used correctly, so such physical memory will be made uncached. Also, note that if this range of physical memory encompasses different types of physical memory that need to be treated differently, this mode will result in different sub-ranges being cached differently, so that the entire range of physical memory is correctly cached. Whenever possible, this mode will result in caching being done to yield the greatest possible efficiency.

CACHE_MODE_INHIBITED - If there is a cache, and if it can be uncached, it is uncached. If this is not possible, but is possible on a sub-range, it is done on any such subranges. Otherwise, nothing is done.

CACHE_MODE_WRITE_THROUGH - If there is a cache, and if it can be converted to write-through mode, it is so converted. If this is not possible, but is possible on a sub-range, it is done on any such sub-ranges. Otherwise, nothing is done.

CACHE_MODE_COPY_BACK - If there is a cache, and if it can be converted to copy-back mode, it is so converted. If this is not possible, but is possible on a sub-range, it is done on any such sub-ranges. Otherwise, nothing is done.

flushCache

[0201]

```
public void flushCache()
```

[0202] Simply flushes the cache if the physical memory is cached. If it is not cached, this method does nothing. Note that the physical caches are flushed for the entire range of addresses of this object.

[0203] The PhysicalMemory class defines the following native methods corresponding to the abstract cache management methods defined in MainMemory. Each of these native methods is implemented by functionality in the microkernel. For more details, refer to the Microkernel Interfaces Chapter.

```
private native void setcachemode (long base, long len, int mode);
```

```
private native int getcachemode(long base, long len);
```

```
private native void flushcache(long base, long len);
```

VirtualAddressSpace

[0204]

```
public final class virtualAddressSpace extends AddressSpace
```

[0205] Each virtual machine has its ownVirtualAddressSpace.

[0206] Besides the constructor, the following public string is defined:

```
public final static String name = "virtual";
```

VirtualAddressSpace

[0207]

```
public VirtualAddressSpace (ExpansionBus b)
```

[0208] This constructor first calls the constructor of the super-class, AddressSpace, specifying the ExpansionBus argument and the String defined above, so that they may be used by the public methods, getExpansionBus() and getName() of AddressSpace. Then, it calls the getAddressSize() and the allocId() native methods (see below) and supplies them to an protected initialization routine of AddressSpace, so that they may then be used by the public methods, getLowestAddress(), getHighestAddress(), and getID() of AddressSpace.

[0209] The VirtualAddressSpace defines the following native methods, implemented in the microkernel, which provide the size of addresses (in bits) and an ID for the address space, respectively. The ID is used when constructing VirtualMemory objects, so that the particular virtual address space for the mapping may be specified to the microkernel.

```
private native int getAddressSize();
```

```
private native int allocId();
```

VirtualMemory

[0210]

```
public abstract class VirtualMemory extends AccessibleMemory
```

[0211] The VirtualMemory object represents memory that can be accessed directly via load and store operations from a CPU, so it follows that VirtualMemory should be a sub-class of AccessibleMemory. The VirtualMemory abstract class does not actually implement any of the access methods that are abstracted by the AccessibleMemory class, as those abstract methods are all finally implemented by sub-classes of VirtualMemory. The VirtualMemory class does provide facilities for managing the mappings of VirtualMemory objects to PhysicalMemory objects, and these facilities are common to all the sub-classes of VirtualMemory.

[0212] VirtualMemory objects are always mapped to some physical memory (which ones it is mapped to cannot be determined unless the mappings are locked). However, even in the case when running on a machine

without an MMU, there are still VirtualMemory objects, and they are simply mapped one to one such that a virtual address is always equal to the underlying physical address.

[0213] The VirtualMemory class is abstract, since it extends AccessibleMemory without implementing any of the abstract methods defined in AccessibleMemory. VirtualMemory does not define any more abstract methods though.

[0214] When constructing a virtualMemory object, note that sometimes it is not possible to map the entire amount of memory specified, so the getMemLength() method of Memory should be called to determine the actual amount of memory mapped.

[0215] The following constructors and public methods are defined:

VirtualMemory

[0216] As the VirtualMemory class is an abstract class, and thus cannot be instantiated, all three of the following protected constructors are used by subclasses. Note that sometimes it is not possible to map the entire amount of memory specified, so the getMemLength() method of Memory should be called to determine the actual amount of memory mapped.

protected VirtualMemory (int vasicd, PhysicalMemory [] PMList, MemoryConstraints constraints) throws InvalidAddressException, AllocationException

Parameters:

vasid - The virtual address space ID as obtained by calling the getID() method of the AddressSpace class, which is the super-class of VirtualAddressSpace.

PMList - An array of PhysicalMemory objects that are to be mapped into virtual memory in the order as they appear in the array.

constraints - The allocation constraints are used now during the allocation and mapping of the virtual memory; the access constraints are stored in the Memory object and used later when it is accessed.

Throws:

InvalidAddressException is thrown if the microkernel is not able to map virtual memory for the specified parameters.

AllocationException is thrown if the microkernel experiences an allocation failure when attempting to allocate virtual memory for the specified mapping

[0217] A request is made to the microkernel to do the mapping, providing the vasicd, arrays of base physical addresses and lengths, allocation constraints, and an indication that the mapping must be locked. The microkernel returns the base virtual address and length

mapped, which this constructor then uses to initialize this Memory object.

protected virtualMemory(int vasicd, Address len, MemoryConstraints constraints) throws InvalidAddressException, AllocationException

Parameters:

vasid - The virtual address space ID as obtained by calling the getID() method of the AddressSpace class, which is the super-class of VirtualAddressSpace.

len - The length of virtual memory to be allocated.

constraints - The allocation constraints are used now during the allocation and mapping of the virtual memory; the access constraints are stored in the Memory object and used later when it is accessed.

Throws:

InvalidAddressException is thrown if the microkernel is not able to map virtual memory for the specified parameters.

AllocationException is thrown if the microkernel experiences an allocation failure when attempting to allocate virtual memory for the specified mapping

[0218] In the case of this second constructor, a request is made to the microkernel to do the mapping, providing the length of memory, and allocation constraints. The microkernel returns the base virtual address and length mapped, which this constructor then uses to initialize this Memory object.

protected VirtualMemory (VirtualMemory superRange, Address offset, Address len) throws InvalidAddressException, AllocationException

Parameters:

superRange - the super-range VirtualMemory object.

offset - The offset within the current VirtualMemory object for the new sub-range.

len - The length of the new sub-range.

Throws:

InvalidAddressException is thrown if either the offset or the offset + newLength is beyond the range of the VirtualMemory object from which the sub-range is being obtained.

AllocationException is thrown if the superrange Memory object is no longer valid.

[0219] This third protected constructor is used by subclasses for their getSubRange() methods. This constructor invokes the corresponding protected constructor of the super-class, AccessibleMemory. Also, a private variable is initialized to indicate that this sub-range VirtualMemory object is not locked. Note that

even if the super-range is locked, we do not initially indicate that the sub-range is locked. This is because we allow the microkernel to maintain a lock count on a particular range of addresses. This means that when we lock this sub-range, the lock count for it would be whatever count already accumulated for super-ranges plus one for this lock.

[0220] Sub-ranges are kept on private linked lists of the super-range VirtualMemory object, so that the unMap() method works as described.

unMap

[0221]

public void unMap() throws AllocationException

Throws:

AllocationException is thrown if the VirtualMemory object's linked list of sub-ranges is non-null, or if there is an internal failure of the unmapping done in the microkernel.

[0222] If the validity of the super-class Memory is set to null (already unmapped), this call returns immediately. This could happen if two threads were trying to call unMap at about the same time, and one got in ahead of the other.

[0223] Provided the VirtualMemory object is not a sub-range, the microkernel is called to do the unmap, specifying a base virtual address and a length. If it is just a sub-range, it is merely removed from its super-range's Vector of sub-ranges. Finally, the validity of the super-class Memory is set to false.

lock

[0224]

public long lock() throws AllocationException

Returns:

The number of bytes of this VirtualMemory object, if it was successfully locked; otherwise returns the number of bytes, starting at the beginning, that could have been locked, had it been a sub-range.

Throws:

AllocationException is thrown if the microkernel is not able to allocate physical memory to back the entire range of virtual memory, or if there is any other failure of the microkernel in getting this memory mapped in.

[0225] If already locked, this method returns immediately. Otherwise, a request is made to the microkernel to get physical memory backing the entire range of virtual memory (if not already done), and to lock the mapping (or to increase the lock count in the case of a sub-range). The microkernel will need to keep track of multi-

ple requests in possibly multiple virtual address spaces that have locked either virtual memory mappings or DMA mappings to physical memory (including a lock count of mappings that are locked to each physical page).

lockContiguous

[0226]

public void lockContiguous() throws AllocationException

Returns:

The number of bytes of this VirtualMemory object if it was successfully locked contiguously; otherwise if it would have been possible to lock a smaller sub-range (starting at the beginning of the object) the total number of bytes that could have been locked contiguously is returned.

Throws:

AllocationException is thrown if the VirtualMemory is already locked, but not contiguously, or if the microkernel is not able to allocate contiguous physical memory to back the entire range of virtual memory, or if there is any other failure of the microkernel in getting this memory mapped in. It should be noted that the microkernel may abort the operation and return a failure, resulting in this exception being thrown, if doing this contiguous locking would cause any locking of sub-ranges to be undone.

[0227] If already locked contiguously, this method returns immediately. Otherwise, provided the VirtualMemory object is not already locked, a request is made to the microkernel to get contiguous physical memory backing the entire range of virtual memory, and to lock the mapping. The microkernel will need to keep track of multiple requests in possibly multiple virtual address spaces that have locked either virtual memory mappings or DMA mappings to physical memory (including a lock count of mappings that are locked to each physical page).

[0228] The reason for not allowing the contiguous locking to be done when the VirtualMemory object is already locked non-contiguously is that there was probably a reason a client locked it non-contiguously, and it is thus better for the software that locked it non-contiguously to do the unlock (and handle any side-effects of that temporary unlocking) rather than causing such side-effects to occur here.

unlock

[0229]

public void unlock() throws AllocationException

Throws:

AllocationException is thrown if there is some failure in the microkernel to unlock the mapping.

[0230] If the VirtualMemory object is already unlocked, this method returns immediately. Note that this causes the microkernel to decrement its lock count for this whole range; it does not affect any sub-ranges that may be locked.

getPhysicalMemory

[0231]

```
public PhysicalMemory [] getPhysicalMemory () 15
throws AllocationException
```

Returns:

An array of PhysicalMemory objects to which this VirtualMemory object is mapped (array order is for increasing virtual address), if this VirtualMemory object is locked (via either lock() or lockContiguous()); otherwise, if this VirtualMemory object is not locked, this method returns null.

Throws:

AllocationException is thrown if there is any failure of the microkernel in obtaining the information regarding the underlying physical memory (array of base addresses and lengths).

[0232] If the VirtualMemory object is not locked, null is returned. Otherwise, if the PhysicalMemory objects have already been obtained for this VirtualMemory object (reference is kept in a private field of the VirtualMemory object), a reference to these pages is simply returned. Otherwise, first a request is made to the microkernel to obtain arrays of base physical addresses and lengths. Then, using the information obtained from the microkernel, each PhysicalMemory object is constructed. Before the array of PhysicalMemory objects is returned, reference to it is kept in a private field of this VirtualMemory object, so that further requests can be handled without involving the microkernel (as long as the VirtualMemory stays locked).

getCacheMode

[0233]

```
public int getcacheMode()
```

Returns:

The cache mode, either MainMemory.CACHE_MODE_INHIBITED, MainMemory.CACHE_MODE_AMBIGUOUS, MainMemory.CACHE_MODE_WRITE_THROUGH, or Main-

Memory.CACHE_MODE_COPY_BACK.

[0234] Note that this method allows a bus manager to determine how the virtual memory is cached if the bus manager did not record how it had specified the virtual memory to be cached, or if CACHE_MODE_DEFAULT had been specified by setCacheMode(), and now the bus manager wants to know how exactly it was cached as default.

setCacheMode

[0235]

```
public void setcacheMode(int mode)
```

Parameters:

mode - Set to one of (see MainMemory):
CACHE_MODE_DEFAULT,
CACHE_MODE_INHIBITED,
CACHE_MODE_WRITE_THROUGH, or
CACHE_MODE_COPY_BACK

[0236] This method allows the bus manager, which should know about cache coherency of various components of the specific platform it supports, to control the caching of virtual memory. For example, on platforms where a virtual memory cache is not coherent with DMA accesses, it would be necessary for the bus manager to temporarily set the cache mode to CACHE_MODE_INHIBITED, while DMA transfers are set up (or to call flushCache() first and somehow guarantee that the CPU(s) will not access the virtual memory while DMA is done).

[0237] Specifically, this method requests the following, depending upon the particular value of mode

CACHE_MODE_DEFAULT - Cache this range of virtual memory according to what the default is for this type of virtual memory, taking into account the caching requirements of the underlying physical memory. Also, note that if this range of virtual memory encompasses different sub-ranges that must be cached differently, this is handled, so that then entire range of virtual memory is correctly cached. Whenever possible, this mode will result in caching being done to yield the greatest possible efficiency.
CACHE_MODE_INHIBITED - If there is a cache, and if it can be uncached, it is uncached. If this is not possible, but is possible on a sub-range, it is done on any such sub-ranges. Otherwise, nothing is done.

CACHE_MODE_WRITE_THROUGH - If there is a cache, and if it can be converted to write-through mode, it is so converted. If this is not possible, but is possible on a sub-range, it is done on any such sub-ranges. Otherwise, nothing is done.

CACHE_MODE_COPY_BACK - If there is a

cache, and if it can be converted to copy-back mode, it is so converted. If this is not possible, but is possible on a sub-range, it is done on any such sub-ranges. Otherwise, nothing is done.

flushCache

[0238]

```
public void flushCache()
```

[0239] Simply flushes the cache if the virtual memory is cached. If it is not cached, this method does nothing. Note that the caches are flushed for the entire range of addresses of this object.

[0240] The VirtualMemory class defines the following native methods. Each of these native methods is implemented by functionality in the microkernel. For more details, refer to the Microkernel Interfaces section.

```
private native int physmap(long[] ret, int vapid,
long[] paddr, long[] sizes, long minaddr, long max-
addr, int align, int cacheMode, boolean locked,
boolean waitformem);
```

```
private native int map(long[] ret, int vapid, long size,
long minaddr, long maxaddr, int align, int cacheM-
ode, boolean locked, boolean waitformem);
```

```
private native int unmap(int vapid, long base, long
len, int howcreated); private native int dolock(long[]
ret, int vapid, long base, long len);
```

```
private native int dolockcontig(long[] ret, int vapid,
long base, long len); private native int downlock(int
vapid, long base, long len);
private native int getphys(int[] ret, int vapid, long[]
pBase, long[] pLen, long base, long len);
```

```
private native void setcachemode(int vapid, long
base, long len, int mode);
```

```
private native int getcachemode(int vapid, long
base, long len);
```

```
private native void flushcache(int vapid, long base,
long len);
```

DMAAddressSpace

[0241]

```
public final class DMAAddressSpace extends
AddressSpace
```

[0242] The DMAAddressSpace is not really used by the memory classes themselves, since it is assumed that a machine only has one DMA address space, but it is provided for completeness, and for use outside of the

memory classes, with a MemoryDescriptor, by various bus managers.

[0243] Besides the constructor, the following public string is defined:

```
public final static String name = "dma";
```

DMAAddressSpace

[0244]

```
public DMAAddressSpace (ExpansionBus b)
```

[0245] This constructor first calls the constructor of the super-class, AddressSpace, specifying the Expansion-Bus argument and the String defined above, so that they may be used by the public methods, getExpansionBus() and getName() of AddressSpace. Then it calls the getAddressSize() and the alloc_id() native methods (see below) and supplies them to an protected initialization routine of AddressSpace, so that they may then be used by the public methods, getLowestAddress(), getHighestAddress() and getID() of AddressSpace.

[0246] The DMAAddressSpace defines the following native methods, implemented in the microkernel, which provide the size of addresses (in bits) and an ID for the address space, respectively.

```
private native int getAddressSize();
```

```
private native int alloc_id();
```

DMAMemory

[0247]

```
Public class DMAMemory extends Memory
```

[0248] The DMAMemory object represents memory that can be accessed only via addressing external to the CPUs on some DMA channel. These DMA addresses are mapped via the DMA hardware to physical memory addresses. Thus, a DMAMemory object is always mapped to PhysicalMemory objects. Since DMAMemory cannot be accessed either directly or indirectly by the CPU(s), DMAMemory is neither a sub-class of AccessibleMemory nor MainMemory. Rather, DMAMemory is a direct sub-class of Memory.

[0249] When constructing a DMAMemory object, note that sometimes it is not possible to map the entire amount of memory specified, so the getMemLength() method of Memory should be called to determine the actual amount of memory mapped.

[0250] The following constructors and public methods are defined:

DMAMemory

[0251]

public DMAMemory (PhysicalMemory [] PMList, MemoryConstraints constraints) throws InvalidAddressException, AllocationException, NotSupported-
Exception

Parameters:

PMList - An array of PhysicalMemory objects that are to be mapped into DMAMemory in the order as they appear in the array.

constraints - The allocation constraints are used now during the allocation and mapping of the virtual memory; the access constraints are ignored, since it is not possible for a CPU to access DMAMemory.

Throws:

NotSupportedException is throw if a call to the microkernel indicates that DMA is not supported on this platform.

InvalidAddressException is thrown if the constructor called for the super-class, Memory, throws InvalidAddressException. if the specified length does not meet the allocation constraints, if the specified PhysicalMemory objects together cannot meet the allocation constraints, or if any of the PhysicalMemory objects is invalid.

AllocationException is thrown if the specified PhysicalMemory objects are not aligned so that they can be mapped to contiguous ranges of the DMAMemory object, or if there is any other internal error of the microkernel that prevents it from making the specified mapping.

[0252] A request is made to the microkernel to allocate the DMA mappings, providing the arrays of base physical addresses and lengths, and the allocation constraints. The microkernel returns the base DMA address and the length of DMA mapped, and then uses this information to initialize the Memory super-class.

public DMAMemory(Address len, MemoryConstraints constraints) throws InvalidAddressException, AllocationException, NotSupportedException

Parameters:

len - The length of DMAMemory to be allocated.

constraints - Only the allocation constraints are used; the access constraints are ignored

Throws:

NotSupportedException is throw if a call to the microkernel indicates that DMA is not supported on this platform.

InvalidAddressException is thrown if the constructor called for the super-class, Memory, throws InvalidAddressException, or if the specified length does not meet the allocation constraints.

AllocationException is thrown if there is any other internal error of the microkernel that prevents it from making the specified mapping.

[0253] A request is made to the microkernel to allocate the DMA mappings, providing the length to allocate and the allocation constraints. The microkernel returns the base DMA address and the length of DMA mapped, and then uses this information to initialize the Memory super-class.

private DMAMemory(DMAMemory superRange, Address offset, Address len) throws InvalidAddressException, AllocationException

Parameters:

superRange - the super-range DMA object. offset - The offset within the current DMA object for the new sub-range.

len - The length of the new sub-range.

Throws:

InvalidAddressException is thrown if either the offset or the offset + newLength is beyond the range of the DMAMemory object from which the sub-range is being obtained, or if the wad sub-range does not meet the allocation constraints of the original DMAMemory object.

AllocationException is thrown if the superrange Memory object is no longer valid.

[0254] This private constructor is used by the getSubRange() method. This constructor invokes the corresponding protected constructor of the super-class, Memory.

toString

[0255]

public String toString()

Returns:

A string representing this DMAMemory object.

unMap

[0256]

public void unMap() throws AllocationException

Throws:

AllocationException is thrown if the DMAMemory object's linked list of sub-ranges is non-null, or if there is an internal failure of the unmapping done in the microkernel.

[0257] If the validity of the super-class Memory is set to null (already unmapped), this call returns immediately. This could happen if two threads were trying to call the unMap() method at about the same time, and one got in ahead of the other. Provided the DMA Memory object is not a sub-range, the microkernel is called to do the DMA deallocation, specifying a base DMA address and a length. If it is just a sub-range, it is merely removed from its superrange's Vector of sub-range. Finally, the validity of the super-class Memory is set to false.

getPhysicalMemory

[0258]

```
public PhysicalMemory[] getPhysicalMemory()
throws AllocationException
```

Returns:

An array of PhysicalMemory objects to which this DMA Memory is mapped (array order is for increasing DMA address).

Throws:

AllocationException is thrown if there is any failure of the microkernel in obtaining the information regarding the underlying physical memory (array of base addresses and lengths).

[0259] If the PhysicalMemory objects have already been obtained for this DMA Memory object (reference is kept in a private field of the DMA Memory object), a reference to these pages is simply returned. Otherwise, first a request is made to the microkernel to obtain arrays of base physical addresses and lengths. Then, using the information obtained from the microkernel, each PhysicalMemory object is constructed. Before the array of PhysicalMemory objects is returned, reference to it is kept in a private field of this DMA Memory object, so that further requests can be handled without involving the microkernel.

getSubRange

[0260]

```
public Memory getSubRange(Address offset,
Address newLength) throws InvalidAddressException
```

Parameters:

offset - The offset within the current DMA Memory object for the new sub-range.

newLength - The length of the new sub-range.

Returns:

A DMA Memory object which is the specified sub-

range.

Throws:

InvalidAddressException is thrown if either the offset or the offset + newLength is beyond the range of the VirtualMemory object from which the sub-range is being obtained.

AllocationException is thrown if the superrange Memory object is no longer valid.

[0261] This method merely calls the private DMA Memory constructor that is used to construct sub-ranges.

toString

[0262]

```
public String toString()
```

Returns:

A string representing this DMA Memory object.

[0263] The DMA Memory class defines the following native methods. Each of these native methods is implemented by functionality in the microkernel.

```
private native boolean supported();
```

```
private native int physmap(long[] ret, long[] paddr,
long[] sizes, long minaddr, long maxaddr, int align,
boolean waitformem);
```

```
private native int map(long[] ret, long size, long
minaddr, long maxaddr, int align, boolean wait-
formem);
```

```
private native int unmap(long base, long len);
```

```
private native int getphys(int[] ret, long[] pBase,
long[] pLen, long base, long len);
```

VirtualIOMemory

[0264]

```
public abstract class VirtualIOMemory extends VirtualMemory
```

[0265] The following constructors and public methods are defined:

VirtualIOMemory

[0266]

```
protected VirtualIOMemory (int vapid, PhysicalMemory[] PMList,
Memoryconstraints constraints) throws InvalidAddressException,
```

AllocationException

protected VirtualIOMemory(int vapid, Address len, MemoryConstraints constraints) throws InvalidAddressException, AllocationException

protected VirtualIOMemory (VirtualIOMemory superrange, Address offset, Address len) throws InvalidAddressException, AllocationException

[0267] Each of these three protected constructors is only used by sub-classes, and merely invoke the corresponding constructor of VirtualMemory. Refer to VirtualMemory for more details.

[0268] The VirtualIOMemory abstract class currently provides no additional functionality but is merely defined in case there are ever common methods between SwappedVirtualIOMemory and UnSwappedVirtualIOMemory. Currently, this class only defines three protected constructors which merely invoke the corresponding constructors of the super-class VirtualMemory.

SwappedVirtualIOMemory

[0269]

public class SwappedVirtualIOMemory extends VirtualIOMemory

[0270] The following constructors and public methods are defined:

SwappedVirtualIOMemory

[0271]

public SwappedVirtualIOMemory (int vapid, PhysicalMemory[] PMList, MemoryConstraints constraints) throws InvalidAddressException, AllocationException, NotSupportedException

public SwappedVirtualIOMemory (int vapid, Address len, MemoryConstraints constraints) throws InvalidAddressException, AllocationException, NotSupportedException

private SwappedVirtualIOMemory (SwappedVirtualIOMemory superrange, Address offset, Address len) throws InvalidAddressException, AllocationException

[0272] As long as SwappedVirtualIOMemory is supported on this platform, each of these three constructors merely invoke the corresponding constructor of VirtualIOMemory, which in turn invokes the corresponding constructor of VirtualMemory. Refer to VirtualMemory for more details. If SwappedVirtualIOMemory is not supported on this platform, NotSupportedException is thrown.

getSubRange

[0273]

public SwappedVirtualIOMemory getSubRange(Address offset, Address newLength) throws InvalidAddressException, AllocationException

Parameters:

offset - The offset within the current SwappedVirtualIOMemory object for the new sub-range.

newLength - The length of the new sub-range.

Returns:

A SwappedVirtualIOMemory object which is the specified sub-range.

Throws:

InvalidAddressException is thrown if either the offset or the offset + newlength is beyond the range of the SwappedVirtualIOMemory object from which the sub-range is being obtained, or if the proposed sub-range does not meet the allocation constraints of the original SwappedVirtualIOMemory object.

AllocationException is thrown if the superrange Memory object is no longer valid.

[0274] This method merely invokes the third constructor method of SwappedVirtualIOMemory, which eventually invokes the third constructor method of VirtualMemory to construct the sub-range.

toString

[0275]

public String toString()

Returns:

A string representing this SwappedVirtualIOMemory object.

[0276] The SwappedVirtualIOMemory class defines the following native methods corresponding to the abstract access methods that are defined in AccessibleMemory. Each of these native methods is implemented by functionality in the microkernel, which must handle both the fact that the memory is byteswapped, and that it is I/O memory which may, depending upon the platform, require additional operations to deal with the fact that I/O memory may not be coherent with other system memory.

protected native boolean supported();

protected native int ndksum(long adr, long len);

protected native void nSetByte(long adr, byte x);

protected native void nSetShort(long adr, short x);

protected native void nSetInt(long adr, int x);

protected native void nSetLong(long adr, long x); 5

protected native void nSetBytes(long adr, byte
bytes[], int bytesOffset, int length);

protected native void nSetIntArray(long adr, byte 10
bytes[], int bytesoffset, int length_int);

protected native byte nGetByte(long adr);

protected native short nGetshort(long adr); 15

protected native int nGetInt(long adr);

protected native long nGetLong(long adr); 20

protected native void nGetBytes(long adr, byte
bytes[], int bytesoffset, int length);

protected native void nGetintArray(long adr, byte 25
bytes[], int bytesOffset, int length_int);

protected native void nRegisterSource(long adr1,
long len);

protected native void nCopy(long adr2, long len, int 30
copyID);

UnSwappedVirtualIOMemory

[0277] 35

public class UnswappedVirtualIOMemory extends
VirtualIOMemory

UnswappedVirtualIOMemory

[0278]

public UnswappedvirtualIOMemory (int vapid,
PhysicalMemory [] PMList, Memoryconstraints 45
constraints) throws InvalidAddressException, Allo-
cationException, NotSupportedException

public UnswappedvirtualIOMemory (int vapid,
Address len,MemoryConstraints constraints) 50
throws InvalidAddressException, AllocationExcep-
tion,NotSupportedException

private UnswappedvirtualIOMemory (Unswapped-
VirtualIOMemory superrange, Address offset, 55
Address len) throws InvalidAddressException, Allo-
cationException

[0279] As long as UnSwappedVirtualIOMemory is supported on this platform, each of these three constructors merely invoke the corresponding constructor of VirtualIOMemory, which in turn invokes the corresponding constructor of VirtualMemory. Refer to VirtualMemory for more details. If UnSwappedVirtualIOMemory is not supported on this platform, NotSupportedException is thrown.

getSubRange

[0280]

public UnswappedvirtualIOMemory getSub-
Range(Address offset, Address newLength) throws
InvalidAddressException, AllocationException

Parameters:

offset - The offset within the current UnswappedVirtualIOMemory object for the new sub-range.
newLength - The length of the new sub-range.

Returns:

A UnSwappedVirtualIOMemory object which is the specified sub-range.

Throws:

InvalidAddressException is thrown if either the offset or the offset + newLength is beyond the range of the UnSwappedVirtualIOMemory object from which the sub-range is being obtained, or if the proposed sub-range does not meet the allocation constraints of the original UnSwappedVirtualIOMemory object. AllocationException is thrown if the superrange Memory object is no longer valid.

[0281] This method merely invokes the third constructor method of UnSwappedVirtualIOMemory which eventually invokes the third constructor method of VirtualMemory to construct the strange.

toString

[0282]

public String toString()

Returns:

A string representing this UnswappedVirtualIOMemory object.

[0283] The UnSwappedVirtualIOMemory class defines the following native methods corresponding to the abstract access methods that are defined in AccessibleMemory. Each of these native methods is implemented by functionality in the microkernel, which must handle the fact that it is I/O memory which may, depending upon the platform, require additional operations to

deal with the fact that I/O memory may not be coherent with other system memory. Unlike the corresponding native methods of UnSwappedVirtualIOMemory the native methods here do not have to deal with byte-swapping.

```
protected native boolean supported();

protected native int ncksum(long adr, long len);

protected native void nSetByte(long adr, byte x);

protected native void nSetShort(long adr, short x);

protected native void nSetInt(long adr, int x);

protected native void nSetLong(long adr, long x);

protected native void nSetBytes(long adr, byte
bytes[], int bytesOffset, int length);

protected native void nSetIntArray(long adr, byte
bytes[], int bytesOffset, int length_int);

protected native byte nGetByte (long adr);

protected native short nGetshort(long adr);

protected native int nGetInt(long adr);

protected native long nGetLong(long adr);

protected native void nGetBytes(long adr, byte
bytes[], int bytesOffset, int length);

protected native void nGetIntArray(long adr, byte
bytes[], int bytesoffset, int length_int);

protected native void nRegisterSource(long adr1,
long len);

protected native void nCopy(long adr2, long len, int
copyID);
```

VirtualRegularMemory

[0284]

```
public abstract class VirtualRegularMemory
extends VirtualMemory
```

[0285] The following constructors and public methods are defined:

VirtualRegularMemory

[0286]

```
5 protected VirtualRegularMemory (int vapid, Physi-
calMemory [] PMList, MemoryConstraints con-
straints) throws InvalidAddressException,
AllocationException

10 protected VirtualRegularMemory(int vapid, Address
len, MemoryConstraints constraints) throws Invali-
dAddressException, AllocationException

15 protected VirtualRegularMemory (VirtualRegu-
larMemory superrange, Address offset, Address
len) throws InvalidAddressException, AllocationEx-
ception
```

[0287] Each of these three protected constructors is only used by sub-classes, and merely invoke the corresponding constructor of VirtualMemory. Refer to VirtualMemory for more details.

[0288] The VirtualRegularMemory abstract class currently provides no additional functionality, but is merely defined in case there are ever common methods between SwappedVirtualRegularMemory and UnswappedVirtualRegularMemory. Currently, this class only defines three protected constructors which merely invoke the corresponding constructors of the super-class VirtualMemory.

swappedvirtualRegularMemory

[0289]

```
35 public class SwappedVirtualRegularMemory
extends VirtualRegularMemory
```

[0290] The following constructors and public methods are defined:

SwappedVirtualRegularMemory

[0291]

```
45 public SwappedVirtualRegularMemory (int vapid,
PhysicalMemory [] PMList, MemoryConstraint con-
straints) throws InvalidAddressException, Alloca-
tionException, NotSupportedException

50 public SwappedVirtualRegularMemory (int vapid,
Address len, MemoryConstraints constraints)
throws InvalidAddressException, AllocationExcep-
tion, NotSupportedException

55 private SwappedVirtualRegularMemory (Swapped-
VirtualRegularMemory superrange, Address offset,
Address len) throws InvalidAddressException, Allo-
```


cationException

[0292] As long as SwappedVirtualRegularMemory is supported on this platform, each of these three constructors merely invoke the corresponding constructor of Virtual Regular Memory, which in turn invokes the corresponding constructor of VirtualMemory. Refer to VirtualMemory for more details. If SwappedVirtualRegularMemory is not supported on this platform, NotSupportedException is thrown.

getSubRange

[0293]

public SwappedVirtualRegularMemory getSubRange (Address offset, Address newLength) throws InvalidAddressException, AllocationException

Parameters:

offset - The offset within the current SwappedVirtualRegularMemory object for the new sub-range.
newLength - The length of the new sub-range.

Returns:

A SwappedVirtualRegularMemory object which is the specified sub-range.

Throws:

InvalidAddressException is thrown if either the offset or the offset + new Length is beyond the range of the SwappedVirtualRegularMemory object from which the sub-range is being obtained, or if the proposed sub-range does not meet the allocation constraints of the original SwappedVirtualRegularMemory object.

AllocationException is thrown if the superrange Memory object is no longer valid.

[0294] This method merely invokes the third constructor method of SwappedVirtualRegularMemory, which eventually invokes the third constructor method of VirtualMemory to construct the sub-range.

toString

[0295]

public String toString()

Returns:

A string representing this SwappedVirtualRegularMemory object.

[0296] The SwappedVirtualRegularMemory class defines the following native methods corresponding to the abstract access methods that are defined in Acces-

sibleMemory. Each of these native methods is implemented by functionality in the microkernel, which must handle both the fact that the memory is byte-swapped.

protected native boolean supported();

protected native int ncksum(long adr, long len);

protected native void nSetByte(long adr, byte x);

protected native void nSetShort(long adr, short x);

protected native void nSetInt(long adr, int x);

protected native void nSetLong(long adr, long x);

protected native void nSetBytes(long adr, byte bytes[], int bytesOffset, int length);

protected native void nSetIntArray(long adr, byte bytes[], int bytesOffset, int length_int);

protected native byte nGetByte(long adr);

protected native short nGetshort(long adr);

protected native int nGetInt(long adr);

protected native long nGetLong(long adr);

protected native void nGetBytes(long adr, byte bytes[], int bytesOffset, int length);

protected native void nGetIntArray(long adr, byte bytes[], int bytesOffset, int length_int);

protected native void nRegisterSource(long adr1, long len);

protected native void nCopy(long adr2, long len, int copyID);

UnSwappedVirtualRegularMemory

[0297]

public class UnSwappedVirtualRegularMemory extends VirtualRegularMemory

[0298] The following constructors and public methods are defined:

UnSwappedVirtualRegularMemory

[0299]

public UnSwappedVirtualRegularMemory (int vapid, PhysicalMemory [] PMList, MemoryCon-

straints constraints) throws `InvalidAddressException`, `AllocationException`, `NotSupportedException`

```
public UnSwappedVirtualRegularMemory (int
    valid, Address len, MemoryConstraints constraints)
    throws InvalidAddressException, AllocationException,
    NotSupportedException
```

```
private UnSwappedVirtualRegularMemory
    (UnSwappedVirtualRegularMemory superrange,
    Address offset, Address len) throws InvalidAddressException,
    AllocationException
```

[0300] As long as `UnSwappedVirtualRegularMemory` is supported on this platform, each of these three constructors merely invoke the corresponding constructor of `VirtualRegularMemory`, which in turn invokes the corresponding constructor of `VirtualMemory`. Refer to `VirtualMemory` for more details. If `UnSwappedVirtualRegularMemory` is not supported on this platform, `NotSupportedException` is thrown.

`getSubRange`

[0301]

```
public UnSwappedVirtualRegularMemory getSubRange
    (Address offset, Address newLength)
    throws InvalidAddressException, AllocationException
```

Parameters:

`offset` - The offset within the current `UnSwappedVirtualRegularMemory` object for the new sub-range.

`newLength` - The length of the new sub-range.

Returns:

A `UnSwappedVirtualRegularMemory` object which is the specified sub-range.

Throws:

`InvalidAddressException` is thrown if either the offset or the offset + `newLength` is beyond the range of the `UnSwappedVirtualRegularMemory` object from which the sub-range is being obtained, or if the proposed sub-range does not meet the allocation constraints of the original `UnSwappedVirtualRegularMemory` object.

`AllocationException` is thrown if the superrange `Memory` object is no longer valid.

[0302] This method merely invokes the third constructor method of `UnSwappedVirtualRegularMemory`, which eventually invokes the third constructor method of `VirtualMemory` to construct the sub-range.

`toString`

[0303]

```
public String toString()
```

Returns:

A string representing this `UnSwappedVirtualRegularMemory` object.

[0304] The `UnSwappedVirtualRegularMemory` class defines the following native methods corresponding to the abstract access methods that are defined in `AccessibleMemory`. Unlike the corresponding native methods of `SwappedVirtualRegularMemory`, the native methods here do not have to deal with byte-swapping.

```
protected native boolean supported();
```

```
protected native int ncksum(long adr, long len);
```

```
protected native void nSetByte(long adr, byte x);
```

```
protected native void nSetShort(long adr, short x);
```

```
protected native void nSetInt(long adr, int x);
```

```
protected native void nSetLong(long adr, long x);
```

```
protected native void nSetBytes(long adr, byte
    bytes[], int bytesOffset, int length);
```

```
protected native void nSetIntArray(long adr, byte
    bytes[], int bytesOffset, int length_int);
```

```
protected native byte nGetByte(long adr);
```

```
protected native short nGetShort(long adr);
```

```
protected native int nGetInt(long adr);
```

```
protected native long nGetLong(long adr);
```

```
protected native void nGetBytes(long adr, byte
    bytes[], int bytesOffset, int length);
```

```
protected native void nGetIntArray(long adr, byte
    bytes[], int bytesOffset, int length_int);
```

```
protected native void nRegisterSource(long adr1,
    long len);
```

```
protected native void nCopy(long adr2, long len, int
    copy ID);
```

PortIOAddressSpace

[0305]

```
public final class PortIOAddressSpace extends 5
    AddressSpace
```

[0306] The PortIOAddressSpace is not really used by the memory classes themselves, since it is assumed that a machine only has one port I/O address space, but it is provided for completeness, and for use outside of the memory classes, with a MemoryDescriptor, by various bus managers.

[0307] Besides the constructor, the following public string is defined:

```
public final static String name="portio";
```

PortIOAddressSpace

[0308]

```
public PortIOAddressSpace (ExpansionBus b)
```

[0309] This constructor first calls the constructor of the super-class, AddressSpace, specifying the ExpansionBus argument and the String defined above, so that they may be used by the public methods, getExpansionBus(), and getName() of AddressSpace. Then it calls the getAddressSize() and the alloc_id() native methods (see below) and supplies them to an protected initialization routine of AddressSpace, so that they may then be used by the public methods, getLowestAddress(), getHighestAddress(), and getID() of AddressSpace.

[0310] The PortIOAddressSpace defines the following native methods, implemented in the microkernel, which provide the size of addresses (in bits) and an ID for the address space, respectively.

```
private native int getAddressSize(); 40
```

```
private native int alloc_id();
```

PortIOMemory

[0311]

```
public abstract class PortIOMemory extends
    AccessibleMemory
```

[0312] The PortIOMemory abstract class provides some minor functionality that is common between SwappedPortIOMemory and UnSwappedPortIOMemory.

[0313] The following constructors and public methods are defined:

PortIOMemory

[0314]

```
protected PortIOMemory (Address p, Address len,
    MemoryConstraints constraints)
```

Parameters:

[0315]

p-BaseAddress**len-length****constraints - Memory Constraints**

[0316] This constructor merely calls the protected constructor of the super-class, AccessibleMemory.

[0317]

```
public void setCacheMode(int mode)
```

Parameters:

[0318]

mode - Requested cache mode.

[0319] This method does nothing; regardless what cache mode is specified, no caching is supported for PortIOMemory, so the mode is always kept as CACHE_MODE_INHIBITED.

getCacheMode

[0320]

```
public int getCacheMode()
```

Returns:
Cache mode.

[0321] This method always returns CACHE_MODE_INHIBITED.

[0322]

```
public void flushCache()
```

[0323] Since there is no cache to flush, this method does nothing.

SwappedPortIOMemory

toString

[0324]

public class SwappedPortIOMemory extends PortIOMemory 5

[0325] The following constructors and public methods are defined:

SwappedPortIOMemory

[0326]

public SwappedPortIOMemory(Address p, Address len, MemoryConstraints constraints) 15

Parameters:

p-BaseAddress

len-length 20

constraints - Memory Constraints

[0327] This constructor merely calls the protected constructor of the super-class, PortIOMemory.

getSubRange

[0328]

protected getSubRange(Address offset, Address newLength) throws InvalidAddressException, AllocationException 30

Parameters:

offset - The offset within the current SwappedPortIOMemory object for the new sub-range. 35

newLength - The length of the new sub-range.

Returns:

A Memory object (actually specifically a Swapped-PortIOMemory object) which is the specified sub-range. 40

Throws:

InvalidAddressException is thrown if either the offset or the offset + newLength is beyond the range of the 45

SwappedPortIOMemory object from which the sub-range is being obtained, or if the proposed sub-range does not meet the constraints (e.g. alignment) of the original SwappedPortIOMemory object. 50

AllocationException is thrown if the superrange Memory object is no longer valid.

[0329] This method merely invokes the constructor method of SwappedPortIOMemory.

[0330]

public String toString()

Returns:

A string representing this SwappedPortIOMemory object.

[0331] The SwappedPortIOMemory class defines the following native methods corresponding to the abstract access methods that are defined in AccessibleMemory. Each of these native methods is implemented by functionality in the microkernel, which must handle both the fact that the memory is byte-swapped, and that it is port memory.

protected native boolean supported();

protected native int ncksum(long adr, long len);

protected native void nSetByte(long adr, byte x);

protected native void nSetShort(long adr, short x);

protected native void nSetInt(long adr, int x);

protected native void nSetLong(long adr, long x);

protected native void nSetBytes(long adr, byte bytes[], int bytesOffset, int length);

protected native void nSetIntArray(long adr, byte bytes[], int bytesOffset, int length_int);

protected native byte nGetByte(long adr);

protected native short nGetShort(long adr);

protected native int nGetInt(long adr);

protected native long nGetLong(long adr);

protected native void nGetBytes(long adr, byte bytes[], int bytesOffset, int length);

protected native void nGetIntArray(long adr, byte bytes[], int bytesOffset, int length_int);

protected native void nRegisterSource(long adr1, long len);

protected native void nCopy(long adr2, long len, int copyID);

UnSwappedPortIOMemory

toString

[0332]

[0338]

public class UnswappedPortIOMemory extends 5
PortIOMemory

public String toString()

[0333] The following constructors and public methods
are defined:

Returns:

A string representing this UnSwappedPortIOMem-
ory object.

UnswappedPortIOMemory

[0334]

[0339] The UnswappedPortIOMemory class defines
the following native methods corresponding to the
abstract access methods that are defined in Accessible-
Memory. Each of these native methods is implemented
by functionality in the microkernel, which must handle
both the fact that the memory is port memory.

public UnSwappedPortIOMemory(Address p, 15
Address len, MemoryConstraints constraints)

Parameters:

p-BaseAddress

len-Length

constraints - Memory Constraints

protected native boolean supported();

protected native int ncksum(long adr, long len);

protected native void nSetByte(long adr, byte x);

protected native void nSetShort(long adr, short x);

protected native void nSetInt(long adr, int x);

[0335] This constructor merely calls the protected
constructor of the super-class, PortIOMemory.

getSubRange

[0336]

protected native void nSetLong(long adr, long x);

protected getSubRange(Address offset, Address 30
newLength) throws InvalidAddressException

protected native void nSetBytes(long adr, byte
bytes[], int bytesOffset, int length);

Parameters:

offset - The offset within the current UnSwapped-
PortIOMemory object for the new sub-range.

newLength - The length of the new sub-range.

protected native void nSetIntArray(long adr, byte
bytes[], int bytesOffset, int length_int);

protected native byte nGetByte(long adr);

Returns:

A Memory object (actually specifically a
UnSwappedPortIOMemory object) which is the
specified sub-range.

protected native short nGetShort(long adr);

protected native int nGetInt(long adr);

protected native long nGetLong(long adr);

Throws:

InvalidAddressException is thrown if either the off-
set or the offset + newLength is beyond the range of
the UnSwappedPortIOMemory object from which
the sub-range is being obtained, or if the proposed
sub-range does not meet the constraints (e.g. align-
ment) of the original UnSwappedPortIOMemory
object.

protected native void nGetBytes(long adr, byte
bytes[], int bytesOffset, int length);

protected native void nGetIntArray(long adr, byte
bytes[], int bytesOffset, int length_int);

protected native void nRegisterSource(long adr1,
long len);

AllocationException is thrown if the superrange
Memory object is no longer valid.

protected native void nCopy(long adr2, long len, int
copyID);

[0337] This method merely invokes the constructor
method of UnswappedPortIOMemory.

Microkernel Memory and DMA Methods

[0340] DMA and Memory methods 80 and 82 of

microkernel 62 consist entirely of the set of native methods of the various Java classes. Together, these native methods are contained in a library of functions for a particular microkernel, and each function of that library either handles the native method itself, or performs a system call to the microkernel to handle the native method.

PhysicalAddressSpace

[0341] The PhysicalAddressSpace class defines the following native methods:

getAddressSize

[0342]

```
private native int getAddressSize();
```

Returns the number of bits of physical addresses

alloc_id

[0343]

```
private native int alloc_id();
```

[0344] Since it is assumed that there is only one physical address space, the microkernel simply returns 0.

PhysicalMemory

[0345] The PhysicalMemory class defines the following native methods corresponding to the abstract cache management methods defined in MainMemory.

[0346] The following native methods control caching of data associated with physical addresses on the particular platform.

setcachemode

[0347]

```
private native void setcachemode(long base, long len, int mode);
```

[0348] This native method sets the caching for physical memory, starting at the physical address, base, and continuing for a length in bytes of len. The caching is set based upon the following possible values of mode (for any other values of mode, do nothing):

- **CACHE_MODE_DEFAULT** - Cache this range of physical memory according to what the default is for this type of physical memory on this platform. Note that, even on platforms that have physical address caches, some physical memory (e.g. I/O registers) may need to be uncached, and it is expected that

the microkernel will make them uncached if they need to be uncached to be used correctly. Also, note that, if the range specified encompasses different types of physical memory that need to be treated differently, the microkernel is expected to cache different sub-ranges differently, if necessary, so that the entire range is cached correctly. Other than these restrictions, the microkernel should try to set up caching to effect the greatest efficiency that is correct. In other words, if caches can be used, while still maintaining data integrity, they should.

- **CACHE_MODE_INHIBITED** - If it is possible to cache this physical memory, and if it is possible to disable its caching, then disable its caching. Otherwise, do nothing. If this is not possible over the entire range, but is possible over one or more sub-ranges, do it at least over any such sub-ranges.
- **CACHE_MODE_WRITE_THROUGH** - If it is possible to cache this physical memory, and if it is possible to set the caching to write through mode, then set it as such. Otherwise, do nothing. If this is not possible over the entire range, but is possible over one or more sub-ranges, do it at least over any such sub-ranges.
- **CACHE_MODE_COPY_BACK** - If it is possible to cache this physical memory and if it is possible to set the caching to copy-back mode, then set it as such. Otherwise, do nothing. If this is not possible over the entire range, but is possible over one or more sub-ranges, do it at least over those sub-ranges.

getcachemode

[0349]

```
private native int getcachemode(long base, long len);
```

[0350] This native method gets the caching for physical memory, starting at the physical address, base, and continuing for a length in bytes of len. One of the following values is to be returned

- **CACHE_MODE_AMBIGUOUS** - If the entire range is not cached the same way, or if none of the following return values apply.
- **CACHE_MODE_INHIBITED** - If there is no cache, or if there is a cache, but this physical memory range is now uncached.
- **CACHE_MODE_WRITE_THROUGH** - If this range of physical memory is now cached in write-through mode.
- **CACHE_MODE_COPY_BACK** - If this range of physical memory is now cached in copy-back mode.

flushcache

[0351]

```
private native void flushcache(long base, long len);
```

[0352] If any sub-range of this range of physical memory is cached in `CACHE_MODE_COPY_BACK` mode, flush any such caches if possible. If it is not possible to flush them, convert them to `CACHE_MODE_INHIBITED`. If even this is not possible, it is assumed that somehow the platform was designed so that DMA is always coherent with respect to any physical caches, and in those cases, it is fine for this method to do nothing.

VirtualAddressSpace

[0353]

The `VirtualAddressSpace` class defines the following native methods:

getAddressSize

[0354]

```
private native int getAddressSize();
```

[0355] Returns the number of bits of virtual addresses for this virtual address space (the one in which the calling thread is running).

alloc_id

[0356]

```
private native int alloc_id();
```

[0357] Each virtual machine runs in a separate virtual address space (only one virtual address space on a machine with just one virtual machine). This method allows the microkernel to specify an ID (only the microkernel will interpret this ID) which can be used subsequently to identify the particular virtual address space for a mapping in native methods of the `VirtualMemory` class. The particular virtual address space is to be the one which the head calling this native method is set up to use.

VirtualMemory

[0358] The `VirtualMemory` class defines the following native methods., which control virtual memory. If a platform does not really have virtual memory, then virtual memory is defined to be mapped (automatically) one-to-one to physical memory in such a way that the virtual address corresponding to a physical address is the

same as that physical address.

[0359] The first three native methods map and unmap virtual memory. The first native method, `physmap()`, sets up virtual memory mapped to an array of physical memory ranges, while the second native method, `map()`, simply maps a specified amount of virtual memory, without specifying the physical memory to which it maps. The third native method, `unmap()`, undoes the mappings that the other two methods create.

physmap

[0360]

```
private native int physmap(long[] ret, int vapid,
long[] paddr, long[] sizes, long minaddr, long max-
addr, int align, int cacheMode, boolean locked,
boolean waitformem);
```

[0361] The particular virtual address space `m` in which this mapping is to be done is specified by `vapid`, which the caller is to have obtained via a call to the `alloc_vapid()` native method of the `VirtualAddressSpace` class.

[0362] The physical memory to be mapped is specified via the arrays `paddr` and `sizes` in such a way that the `n`th element of `paddr` and the `n`th element of `sizes` specify the base physical address and length for a range of physical memory, respectively. These ranges of physical memory are to be mapped to increasingly higher virtual addresses in such a way that physical memory range `n` (as specified by `n`th elements of `paddr` and `sizes`) is to be mapped to virtual addresses which are lower than the virtual addresses which map to physical memory range `n+1`. Furthermore, the virtual address ranges are to be contiguous, so that they form one virtual address range.

[0363] For this mapping, any range of virtual addresses can be used, except that all virtual addresses in the range must be greater than or equal to `minaddr`, and all virtual addresses in the range must be less than or equal to `maxaddr`. Also, the base virtual address is to be aligned as specified by `align`, which could have one of the following values:

- `MemoryConstraints.ALIGN_BYTE` (0) - no alignment restriction.
- `MemoryConstraints.ALIGN_SHORT` (1) - align to short.
- `MemoryConstraints.ALIGN_INT` (2) - align to int.
- `MemoryConstraints.ALIGN_LONG` (3) - align to long.
- `MemoryConstraints.ALIGN_CACHE` (4) - align to cache line size.
- `MemoryConstraints.ALIGN_PAGE` (5) - align to page size.

[0364] The virtual memory is to be cached as though `cacheMode` were specified for mode in a call to the

native method, setcachemode().

[0365] If locked is true, the virtual memory mappings to the particular physical pages are to be locked as though the dolock() native method had been called.

[0366] If waitformem is true and it is currently not possible to complete the mapping request, due to a resource running out temporarily, then wait indefinitely to complete the request. Otherwise, do not wait.

[0367] If not all these constraints can be met, but they could be met for a smaller sub-range, which must start with the first range of physical memory, then do the mapping for that smaller sub-range. For example, if this platform does not really have virtual memory, and if the physical memory ranges are discontinuous, then obviously only the first range of physical memory can be mapped.

[0368] The base virtual address of the range is to be returned in ret[0] and the length of virtual memory mapped is to be returned in ret[1].

[0369] The possible return values are:

- SUCCESS - The mapping was successful. Either the entire amount of physical memory was mapped, or at least a sub-range was mapped.
- ADDRESS_FAILURE - There was some failure related to the physical addresses specified. For whatever reason those addresses are undefined physical memory or physical memory for which mappings can not be setup.
- ALLOC_FAILURE - There was some other fatal failure in the attempt to create the mappings.

map

[0370]

```
private native int map(long[] ret, int vapid, long size,
    long minaddr, long maxaddr, int align, int cacheM-
    ode, boolean locked, boolean waitformem);
```

[0371] The particular virtual address space in which this mapping is to be done is specified by vapid, which the caller is to have obtained via a call to the alloc_vapid() native method of the VirtualAddressSpace class.

[0372] The amount of virtual memory to allocate is specified by size in bytes. For this mapping, any range of virtual addresses can be used, except that all virtual addresses in the range must be greater than or equal to minaddr, and all virtual addresses in the range must be less than or equal to maxaddr. Also, the base virtual address is to be aligned as specified by align, which could have one of the following values:

- MemoryConstraints.ALIGN_BYTE (0) - no alignment restriction.
- MemoryConstraints.ALIGN_SHORT (1) - align to short.

- MemoryConstraints.ALIGN_INT (2) - align to int.
- MemoryConstraints.ALIGN_LONG (3) - align to long.
- MemoryConstraints.ALIGN_CACHE (4) - align to cache line size.
- MemoryConstraints.ALIGN_PAGE (5) - align to page size.

[0373] The virtual memory is to be cached as though cacheMode were specified for mode in a call to the native method, setcachemode().

[0374] If locked is true, the virtual memory mappings to the particular physical pages are to be locked as though the dolock() native method had been called.

[0375] If waitformem is true and it is currently not possible to complete the mapping request, due to a resource running out temporarily, then wait indefinitely to complete the request. Otherwise, do not wait.

[0376] If not all these constraints can be met, but they could be met for a smaller sub-range, which must start with the first range of physical memory; then do the mapping for that smaller sub-range.

[0377] The base virtual address of the range is to be returned in ret[0] and the length of virtual memory mapped is to be returned in ret[1].

[0378] The possible return values are:

- SUCCESS - The mapping was successful. Either the entire amount of physical memory was mapped, or at least a sub-range was mapped.
- ADDRESS_FAILURE - There was some failure related to the physical addresses specified. For whatever reason those addresses are undefined physical memory or physical memory for which mappings can not be set up.
- ALLOC_FAILURE - There was some other fatal failure in the attempt to create the mappings.

unmap

[0379]

```
private native int unmap(int vapid, long base, long
    len, int howcreated);
```

[0380] The particular virtual address space in which this unmapping is to be done is specified by vapid, which the caller is to have obtained via a call to the alloc_vapid() native method of the VirtualAddressSpace class.

[0381] The virtual memory to be unmapped is that starting at base and having a length of len in bytes.

[0382] The howcreated parameter specifies how the virtual memory was mapped in the first place:

- CREATED_MAP - Created via physmap().
- CREATED_MALLOC - Created via map().

[0383] The possible return values are:

- SUCCESS - The unmapping was a success. The entire range specified was unmapped.
- ALLOC_FAILURE - There was some other fatal failure in the attempt to do the unmapping.

[0384] The next three native methods lock and unlock virtual memory mappings. The native method, `dolock()`, locks virtual memory to particular physical memory without regard to the underlying physical memory being contiguous, while the native method, `dolockcontig()`, does the locking to contiguous physical memory. The native method, `unlock()`, unlocks the mappings that the other two methods lock.

`dolock`

[0385]

```
private native int dolock(long[] ret, int vaside, long
base, long len);
```

[0386] The particular virtual address space in which this locking is to be done is specified by `vaside`, which the caller is to have obtained via a call to the `alloc_vaside()` native method of the `VirtualAddressSpace` class.

[0387] The virtual memory to be locked is that starting at `base` and having a length of `len` in bytes.

[0388] When a range of virtual memory is locked, it means that the microkernel cannot change the identity of the underlying physical memory. Thus, while locked, a particular range of virtual memory always maps to the same range of physical memory. It should be noted that the microkernel needs to maintain a lock count, such that each call to lock a particular range must be undone by as many calls to the `dounlock()` native method.

[0389] If the locking is not successful, this native method will not cause any locking of a partial range to be done. If the locking was successful, `ret[0]` will be the total number of bytes of the range that are locked (the entire length of the range). If the locking was not successful, but if it would have been possible to lock a sub-range starting at the beginning of the range, then `ret[0]` is the length of that sub-range.

[0390] The possible return values are:

- SUCCESS - The locking was a success. The entire range specified was locked.
- NOT_ALL_LOCKED - it would not be possible to lock the entire specified range, but it would be possible to lock a smaller sub-range starting at the beginning, whose length is `ret[0]`.
- ALLOC_FAILURE - There was some other fatal failure in the attempt to do the unmapping.

`dolockcontig`

[0391]

```
private native int dolockcontig(long[] ret, int vaside,
long base, long len);
```

[0392] The particular virtual address space in which this locking is to be done is specified by `vaside`, which the caller is to have obtained via a call to the `alloc_vaside()` native method of the `VirtualAddressSpace` class.

[0393] The virtual memory to be locked is that starting at `base` and having a length of `len` in bytes.

[0394] When a range of virtual memory is locked, it means that the microkernel cannot change the identity of the underlying physical memory. Thus, while locked, a particular range of virtual memory always maps to the same range of physical memory. It should be noted that the microkernel needs to maintain a lock count, such that each call to lock a particular range must be undone by as many calls to the `dounlock()` native method.

[0395] As opposed to the `dolock()` native method, this particular native method is to do the locking so that the entire range of virtual addresses map to contiguous physical memory. In order to do this, the microkernel may need to remap portions of this virtual range. Note, though, that if this remapping violates any locking already done via `dolock()` or `dolockcontig()`, then this `dolockcontig()` call is to be failed (with NOT_ALL_LOCKED).

[0396] If the locking is not successful, this native method will not cause any locking of a partial range to be done. If the locking was successful, `ret[0]` will be the total number of bytes of the range that are locked (the entire length of the range). If the locking was not successful, but if it would have been possible to lock a sub-range starting at the beginning of the range, then `ret[0]` is the length of that sub-range.

[0397] The possible return values are:

- SUCCESS - The locking was a success. The entire range specified was locked contiguously.
- NOT_ALL_LOCKED - it would not be possible to lock the entire specified range, but it would be possible to lock contiguously a smaller sub-range starting at the beginning, whose length is `ret[0]`.
- ALLOC_FAILURE - There was some other fatal failure in the attempt to do the unmapping.

`dounlock`

[0398]

```
private native int dounlock(int vaside, long base, long
len);
```

[0399] The particular virtual address space in which this unlocking is to be done is specified by `vaside`, which

the caller is to have obtained via a call to the `alloc_vasid()` native method of the `VirtualAddressSpace` class.

[0400] The virtual memory to be unlocked is that starting at base and having a length of len in bytes. This call decrements the lock count for the range which had been incremented by either the `dolock()` or the `dolockcontig()` native methods.

[0401] The possible return values are:

- SUCCESS - The unmapping was a success. The entire range specified was unlocked.
- ALLOC_FAILURE - There was some other fatal failure in the attempt to do the unlocking.

getphys

[0402]

```
private native int getphys(int[] ret, int vasid, long[]
pBase, long[] pLen, long base, long len);
```

[0403] Provided the virtual memory is first locked, it is guaranteed that the identity of the underlying physical memory will remain the same. A call to the `getphys()` native method is used to obtain the underlying physical memory.

[0404] The particular virtual address space in which this operation is to be done is specified by `vasid`, which the caller is to have obtained via a call to the `alloc_vasid()` native method of the `VirtualAddressSpace` class.

[0405] The base of the range of virtual memory is specified by `base`, and the length of the range by `len` in bytes.

[0406] If the virtual memory is not locked, this call does nothing and simply returns FAILURE.

[0407] The physical memory is to be specified in the supplied arrays `pBase` and `pLen`, such that the `pBase[0]` and `pLen[0]` are the base and length of the physical memory that starts mapping the beginning of the virtual memory, and such that `(pBase[n+1], pLen[n+1])` specifies physical memory to which a higher range of virtual addresses is mapped than the range of virtual addresses that map the physical memory range specified by `(pBase[n], pLen[n])`.

[0408] It is not possible for the caller to know how many discontinuous ranges of physical memory map this virtual memory range, so it is possible that the supplied arrays for `pBase` and `pLen` may be too small. If this is the case, the microkernel is to return `ARRAY_TOO_SMALL` and specify the number of elements needed in the array in `ret[0]`.

[0409] If this call is successful, the number of bits of the physical memory on this platform is specified in `ret[1]`.

[0410] The next native methods control caching of data associated with virtual addresses on the particular

platform.

setcachemode

[0411]

```
private native void setcachemode(int vasid, long
base, long len, int mode);
```

[0412] The particular virtual address space in which this operation is to be done is specified by `vasid`, which the caller is to have obtained via a call to the `alloc_vasid()` native method of the `VirtualAddressSpace` class.

[0413] This native method sets the caching for virtual memory, starting at the virtual address, `base`, and continuing for a length in bytes of `len`. The caching is set based upon the following possible values of `mode` (for any other values of `mode`, do nothing):

- `CACHE_MODE_DEFAULT` - Cache this range of virtual memory according to what the default is for this type of virtual memory on this platform. Note that, even on platforms that have virtual address caches, depending upon the underlying physical memory mapped, the virtual memory mapping some physical memory (e.g. I/O-registers) may need to be uncached, and it is expected that the microkernel will make these virtual memory mappings uncached if they need to be uncached to be used correctly. Also, note that if the range specified encompasses virtual memory mapped to different types of physical memory that need to be treated differently, the microkernel is expected to cache different sub-ranges differently, if necessary, so that the entire range is cached correctly. Other than these restrictions, the microkernel should try to set up caching to effect the greatest efficiency that is correct. In other words, if caches can be used, while still maintaining data integrity, they should.
- `CACHE_MODE_INHIBITED` - If it is possible to cache this virtual memory, and if it is possible to disable its caching, then disable its caching. Otherwise, do nothing. If this is not possible over the entire range, but is possible over one or more sub-ranges, do it at least over any such sub-ranges.
- `CACHE_MODE_WRITE_THROUGH` - If it is possible to cache this virtual memory, and if it is possible to set the cache to write-through mode, then set it as such. Otherwise, do nothing. If this is not possible over the entire range, but is possible over one or more sub-ranges, do it at least over any such sub-ranges.
- `CACHE_MODE_COPY_BACK` - If it is possible to cache this virtual memory, and if it is possible to set the caching to copy-back mode, then set it as such. Otherwise, do nothing. If this is not possible over the entire range, but is possible over one or more

sub-ranges, do it at least over those sub-ranges.

getcachemode

[0414]

```
private native int getcachemode(int vapid, long
base, long len);
```

[0415] This native method gets the caching for virtual memory, starting at the virtual address, base, and continuing for a length in bytes of len. One of the following values is to be returned:

- **CACHE_MODE_AMBIGUOUS** - If the entire range is not cached the same way, or if none of the following return values apply.
- **CACHE_MODE_INHIBITED** - If there is no cache, or if there is a cache, but this virtual memory range is now uncached.
- **CACHE_MODE_WRITE_THROUGH** - If this range of virtual memory is now cached in write-through mode.
- **CACHE_MODE_COPY_BACK** - If this range of virtual memory is now cached in copy-back mode.

flushcache

[0416]

```
private native void flushcache(int vapid, long base,
long len);
```

[0417] If any sub-range of this range of virtual memory is cached in **CACHE_MODE_COPY_BACK** mode, flush any such caches if possible. If it is not possible to flush them, convert them to **CACHE_MODE_INHIBITED**.

DMAAddressSpace

[0418] The **DMAAddressSpace** class defines the following native methods:

getAddressSize

[0419]

```
private native int getAddressSize();
```

Returns the number of bits of DMA addresses

alloc_id

[0420]

```
private native int alloc_id();
```

[0421] Since it is assumed that there is only one DMA

address space, the microkernel simply returns 0.

DMAMemory

5 [0422] The **DMAMemory** class defines the following native methods. Each of these native methods is implemented by functionality in the microkernel.

supported

10 [0423]

```
private native boolean supported();
```

[0424] Returns true if DMA is supported on this platform; otherwise returns false.

physmap

[0425]

```
private native int physmap(long[] ret, long[] paddr,
long[] sizes, long minaddr, long maxaddr, int align,
boolean waitformem);
```

25 [0426] The physical memory to be mapped is specified via the arrays **paddr** and **sizes** in such a way that the *n*th element of **paddr** and the *n*th element of **sizes** specify the base physical address and length for a range of physical memory, respectively. These ranges of physical memory are to be mapped to increasingly higher DMA addresses in such a way that physical memory range *n* (as specified by *n*th elements of **paddr** and **sizes**) is to be mapped to DMA addresses which are lower than the DMA addresses which map to physical memory range *n*+1. Furthermore, the DMA address ranges are to be contiguous, so that they form one DMA address range.

[0427] For this mapping, any range of DMA addresses can be used, except that all DMA addresses in the range must be greater than or equal to **minaddr**, and all DMA addresses in the range must be less than or equal to **maxaddr**. Also, the base DMA address is to be aligned as specified by **align**, which could have one of the following values:

- 45 • **MemoryConstraints.ALIGN_BYTE** (0) - no alignment restriction.
- **MemoryConstraints.ALIGN_SHORT** (1) - align to short.
- 50 • **MemoryConstraints.ALIGN_INT** (2) - align to int.
- **MemoryConstraints.ALIGN_LONG** (3) - align to long.
- **MemoryConstraints.ALIGN_CACHE** (4) - align to cache line size.
- 55 • **MemoryConstraints.ALIGN_PAGE** (5) - align to page size.

[0428] If **waitformem** is true and it is currently not pos-

sible to complete the mapping request, due to a resource running out temporarily, then wait indefinitely to complete the request. Otherwise, do not wait.

[0429] If not all these constraints can be met, but they could be met for a smaller sub-range, which must start with the first range of physical memory, then do the mapping for that smaller sub-range. For example, if this platform does not really have DVMA (Direct Virtual Memory Addressing), and if the physical memory ranges are discontiguous, then obviously only the first range of physical memory can be mapped.

[0430] The base DMA address of the range is to be returned in ret[0] and the length of DMA memory mapped is to be returned in ret[1].

[0431] The possible return values are:

- SUCCESS - The mapping was successful. Either the entire amount of physical memory was mapped, or at least a sub-range was mapped.
- ADDRESS_FAILURE - There was some failure related to the physical addresses specified. For whatever reason those addresses are undefined physical memory or physical memory for which mappings can not be set up.
- ALLOC_FAILURE - There was some other fatal failure in the attempt to create the mappings.

map

[0432]

```
private native int map(long[] ret, long size, long
minaddr, long maxaddr, int align, boolean wait-
formem);
```

[0433] The amount of DMA memory to allocate is specified by size in bytes. For this mapping, any range of DMA addresses can be used, except that all DMA addresses in the range must be greater than or equal to minaddr, and all DMA addresses in the range must be less than or equal to maxaddr. Also, the base DMA address is to be aligned as specified by align, which could have one of the following values:

- MemoryConstraints.ALIGN_BYTE (0) - no alignment restriction.
- MemoryConstraints.ALIGN_SHORT (1) - align to short.
- MemoryConstraints.ALIGN_INT (2) - align to int.
- MemoryConstraints.ALIGN_LONG (3) - align to long.
- MemoryConstraints.ALIGN_CACHE (4) - align to cache line size.
- MemoryConstraints.ALIGN_PAGE (5) - align to page size.

[0434] If waitformem is true and it is currently not possible to complete the mapping request, due to a

resource running out temporarily, then wait indefinitely to complete the request. Otherwise, do not wait.

[0435] If not all these constraints can be met, but they could be met for a smaller sub-range, which must start with the first range of physical memory, then do the mapping for that smaller sub-range.

[0436] The base DMA address of the range is to be returned in ret[0] and the length of DMA memory mapped is to be returned in ret[1]. The possible return values are:

- SUCCESS - The mapping was successful. Either the entire amount of physical memory was mapped, or at least a sub-range was mapped.
- ADDRESS_FAILURE - There was some failure related to the physical addresses specified. For whatever reason those addresses are undefined physical memory or physical memory for which mappings can not be set up.
- ALLOC_FAILURE - There was some other fatal failure in the attempt to create the mappings.

unmap

[0437]

```
private native int unmap(long base, long len);
```

[0438] The DMA memory to be unmapped is that starting at base and having a length of len in bytes.

[0439] The possible return values are:

- SUCCESS - The unmapping was a success. The entire range specified was unmapped.
- ALLOC_FAILURE - There was some other fatal failure in the attempt to do the unmapping.

getphys

[0440]

```
private native int getphys(int[] ret, long[] pBase,
long[] pLen, long base, long len);
```

[0441] The base of the range of DMA memory is specified by base, and the length of the range by len in bytes.

[0442] The physical memory is to be seed in the supplied arrays pBase and pLen, such that the pBase[0] and pLen[0] are the base and length of the physical memory that starts mapping the beginning of the DMA memory, and such that (pBase[n+1], pLen[n+1]) specifies physical memory to which a higher range of DMA addresses is mapped than the range of DMA addresses that map the physical memory range specified by (pBase[n], pLen[n]).

[0443] It is not possible for the caller to know how many discontiguous ranges of physical memory map this DMA memory range, so it is possible that the sup-

plied arrays for pBase and pLen may be too small. If this is the case, the microkernel is to return ARRAY_TOO_SMALL and specify the number of elements needed in the array in ret[0].

[0444] If this call is successful, the number of bits of the physical memory on this platform is specified in ret[1].

AccessibleMemory Sub-Classes

[0445] The various non-abstract sub-classes of AccessibleMemory define as native methods the methods abstracted in AccessibleMemory. These sub-classes consist of SwappedVirtualIOMemory, UnSwappedVirtualIOMemory, SwappedVirtualRegularMemory, UnSwappedVirtualRegularMemory, SwappedPortIOMemory, and UnSwappedPortIOMemory. For each of these sub-classes there are separate versions of these native methods.

[0446] These separate versions of native methods are mostly similar to one another, with their differences specified by the following flavors:

- Swapped/UnSwapped - Versions of the native methods for swapped classes (SwappedVirtualIOMemory, SwappedVirtualRegularMemory, and SwappedPortIOMemory) must byte-swap (reverse the endianness) any data they get from the native methods before operating upon it, and must byte-swap back any data they supply to the classes. The other classes (UnSwappedVirtualIOMemory, UnSwappedVirtualRegularMemory, and UnSwappedPortIOMemory) are not to have any such byte-swapping done.
- IO/Regular - Classes supporting IO memory (SwappedVirtualIOMemory and UnSwappedVirtualIOMemory) may require special platform-specific (implementation dependent) operations to be done whenever accesses are done to that memory. Classes supporting regular RAM memory (SwappedVirtualRegularMemory and UnSwappedVirtualRegularMemory) do not require any such special platform-specific operations. Note that on some platforms, there is no difference.
- VirtualPort - Classes for virtual memory (SwappedVirtualIOMemory, UnSwappedVirtualIOMemory, SwappedVirtualRegularMemory and UnSwappedVirtualRegularMemory) can assume that memory is accessed via normal load and store instructions, and that the addresses specified in the native methods refer to virtual memory addresses. Classes for port memory (SwappedPortIOMemory and UnSwappedPortIOMemory) must do special I/O instructions to access memory.

[0447] Modulo the above specification of flavors, the following is specified for each of these six sets of native methods:

supported

[0448]

protected native boolean supported();

[0449] Returns true if this particular sub-class of AccessibleMemory is supported on this platform; otherwise returns false.

ncksum

[0450]

protected native int ncksum(long adr, long len);

[0451] Computes a sum of the bytes of the memory (modulo 0x 10000) starting at the address adr for the length len in bytes.

nSetByte

[0452]

protected native void nsetByte(long adr, byte x);

[0453] This operation must be done with respect to the following particular flavors (see above) of the specific subclass of AccessibleMemory which this native method supports: IO/Regular, Virtual/Port.

[0454] The byte value, x, is to be written at the address, adr.

nSetShort

[0455]

protected native void nsetshort(long adr, short x);

[0456] This operation must be done with respect to the following particular flavors (see above) of the specific sub-class of AccessibleMemory which this native method supports: Swapped/UnSwapped, IO/Regular, Virtual/Port.

[0457] The short value, x, is to be written at the address, adr.

nSetInt

[0458]

protected native void nsetInt(long adr, int x);

[0459] This operation must be done with respect to the following particular flavors (see above) of the specific sub-class of AccessibleMemory which this native method supports: Swapped/UnSwapped, IO/Regular, Virtual/Port.

[0460] The int value, x, is to be written at the address, adr.

nSetLong

[0461]

protected native void nsetlong(long adr, long x);

[0462] This operation must be done with respect to the following particular flavors (see above) of the specific sub-class of AccessibleMemory which this native method supports: Swapped/UnSwapped, IO/Regular, Virtual/Port.

[0463] The long value, x, is to be written at the address, adr.

nSetBytes

[0464]

protected native void nsetBytes(long adr, byte bytes[], int bytesOffset, int length);

[0465] This operation must be done with respect to the following particular flavors (see above) of the specific sub-class of AccessibleMemory which this native method supports: Swapped/UnSwapped, IO/Regular, Virtual/Port.

[0466] Starting at the offset, bytesOffset, of the array, bytes, the array of bytes is to be written starting at the address adr, for a total length in bytes of length.

nSetIntArray

[0467]

protected native void nsetIntArray(long adr, byte bytes[], int bytesoffset, int length_int);

[0468] This operation must be done with respect to the following particular flavors (see above) of the specific sub-class of AccessibleMemory which this native method supports: Swapped/UnSwapped, IO/Regular, Virtual/Port.

[0469] Starting at the offset, bytesOffset, of the array, bytes, the array of bytes is to be written, treating it as an array of integers, starting at the address adr, for a total length in integers of length_int.

nGetByte

[0470]

protected native byte nGetByte(long adr);

[0471] This operation must be done with respect to the following particular flavors (see above) of the specific sub-class of AccessibleMemory which this native

method supports: IO/Regular, Virtual/Port.

[0472] Read and return a byte from the address, adr.

nGetShort

[0473]

protected native short nGetshort(long adr);

[0474] This operation must be done with respect to the following particular flavors (see above) of the specific sub-class of AccessibleMemory which this native method supports: Swapped/UnSwapped, IO/Regular, Virtual/Port.

[0475] Read and return a short from the address, adr.

nGetInt

[0476]

protected native int nGetInt(long adr);

[0477] This operation must be done with respect to the following particular flavors (see above) of the specific sub-class of AccessibleMemory which this native method supports: Swapped/UnSwapped, IO/Regular, Virtual/Port.

[0478] Read and return an int from the address, adr.

nGetLong

[0479]

protected native long nGetLong(long adr);

[0480] This operation must be done with respect to the following particular flavors (see above) of the specific sub-class of AccessibleMemory which this native method supports: Swapped/UnSwapped, IO/Regular, Virtual/Port.

[0481] Read and return a long from the address, adr.

nGetBytes

[0482]

protected native void nGetBytes(long adr, byte bytes[], int bytesOffset, int length);

[0483] This operation must be done with respect to the following particular flavors (see above) of the specific sub-class of AccessibleMemory which this native method supports: Swapped/UnSwapped, IO/Regular, Virtual/Port.

[0484] Start reading bytes from adr for a total of length bytes, and read them into the array bytes, starting at the offset of bytesOffset of that array.

nGetIntArray**[0485]**

protected native void nGetIntArray(long adr, byte 5
bytes[], int bytesOffset, int length_int);

[0486] This operation must be done with respect to the following particular flavors (see above) of the specific sub-class of AccessibleMemory which this native method supports: Swapped/UnSwapped, IO/Regular, Virtual/Port. 10

[0487] Start reading ints from adr for a total of length_int ints, and read them into the array bytes, starting at the offset of bytesOffset of that array. 15

nRegisterSource**[0488]**

protected native void nRegisterSource(long adr1, 20
long len);

[0489] This operation registers the source of a copy operation, and returns an ID, which can later be passed to the nCopy native method. The base address of the source, adr1, and the length in bytes to copy, len, are to be registered. It is also necessary to take into account and register the following flavors depending upon the particular nature of the source memory: 25
Swapped/UnSwapped, IO/Regular, Virtual/Port. 30

nCopy**[0490]**

protected native void nCopy(long adr2, long len, int 35
copyID);

[0491] This operation performs a copy from the source as specified by copyID, to the destination specified here. 40

[0492] If copyID does not correspond to a registered source, do nothing.

[0493] The destination address, adr2, and length in bytes, len, are as specified. The source address and length may be obtained by what was registered for the source, as specified by copyID. If the length of source and destination do not agree, use the lesser of the two. 45

[0494] In performing this copy, it is necessary to take into account the following flavors for both the source and destination: Swapped/UnSwapped, IO/Regular, Virtual/Port. 50

[0495] The classes described above are stored on a computer readable medium, such as floppy disks, a CD-ROM, or optical disk. Alternatively, they are supplied to computer system 10 in a read-only memory, or provided in the form of a computer data carrier wave over a network. 55

[0496] It will be apparent to those skilled in the art that various modifications and variations can be made in the disclosed process and product without departing from the scope or spirit of the invention. Other embodiments of the invention will be apparent to those skilled in the art from consideration of the specification and practice of the invention disclosed herein. It is intended that the specification and examples be considered as exemplary only, with a true scope and spirit of the invention being indicated by the following claims.

[0497] The features disclosed in the foregoing description, in the claims and/or in the accompanying drawings may, both separately and in any combination thereof, be material for realising the invention in diverse forms thereof.

Claims

1. A method for providing memory functionality to object-oriented client software components of a computer system having a CPU, comprising the steps of: 20

providing a first set of memory classes, each class of the first set being platform-independent;

providing a second set of memory classes, each class of the second set being a subclass of a class of the first set and being platform-dependent; and

performing client component memory functions by accessing only objects of classes of the first set. 35

2. A method as recited in claim 1, wherein the step of providing a first set of memory classes includes the substep of providing a first abstract memory class having only the attributes of base address, length, and constraints.

3. A method as recited in claim 2, wherein the step of providing a first set of memory classes includes the substep of providing a first instantiable class of memory which extends the first abstract memory class and which provides no functionality beyond that provided by the first abstract memory class.

4. A method as recited in claim 2, wherein the step of providing a first set of memory classes includes the substep of providing a second abstract memory class which extends the first abstract memory class and which specifies only abstract methods for controlling cache memory. 50

5. A method as recited in claim 4, wherein the step of providing a first set of memory classes includes the 55

substep of providing a third abstract memory class extending the second abstract memory class and representing all memory that can be directly accessed by the CPU.

6. A method as recited in claim 5, wherein the substep of providing a third abstract memory class includes the substep of providing a third abstract memory class which defines platform-independent abstract methods for accessing memory.
7. A method as recited in claim 1, wherein the step of providing a second set of memory classes includes the substep of providing a second instantiable memory class, extending the second abstract memory class, which represents physical memory.
8. A method as recited in claim 7, wherein the substep of providing a second instantiable memory class comprises the substep of providing a second instantiable memory class which provides implementations of cache management methods abstracted in the second abstract memory class.
9. A method as recited in claim 8, wherein the step of providing a second set of memory classes includes the substep of providing a fourth abstract memory class, which extends the third abstract memory class, and which contains methods for setting up mappings to objects of the second instantiable memory class and a method for allocating a specific amount of virtual memory without specifying a particular physical memory.
10. A method as recited in claim 1, wherein the step of providing a second set of memory classes includes the substep of providing a third instantiable memory class, which extends the first abstract memory class, for setting up DMA mapping to physical memory and for doing corresponding unmapping.
11. A method as recited in claim 1, wherein the step of providing a second set of memory classes includes the substep of providing a second set of memory classes in which methods are separated into two groups, the first group being written in non-native language and the second group being written in native language.
12. A method as recited in claim 11, wherein the step of providing a second set of memory classes includes the substep of providing the second group of methods in a library of functions which access the micro-kernel.
13. A method as recited in claim 12, wherein:

the step of providing a first set of memory

classes includes the substep of providing a second abstract memory class which extends the first abstract memory class and which specifies only abstract methods for caching memory; and

the step of providing a second set of memory classes includes the substeps of providing native methods in the second group of methods corresponding to the abstract cache management methods defined in the second abstract memory class.

14. Apparatus for performing memory functions in an object-oriented computer system, comprising:

a first set of memory classes, each class of the first set being platform-independent;

a second set of memory classes, each class of the second set being a subclass of a class of the first set and being platform-dependent;

a device driver accessing only objects of classes of the first set; and

a bus manager accessing classes of the second set.

15. A computer-readable medium containing instructions for providing memory functionality to object-oriented client software components of a computer system having a CPU, comprising:

a first set of memory classes, each class of the first set being platform-independent; and

a second set of memory classes, each class of the second set being a subclass of a class of the first set and being platform-dependent.

16. A computer-readable medium as in 15, wherein the second set of classes includes a first subset of methods written in a platform-independent language and a second subset of methods written in native code.

17. A computer data signal on a carrier wave containing instructions for providing memory functionality to object-oriented client software components of a computer system having a CPU, comprising:

a first set of memory classes, each class of the first set being platform-independent; and

a second set of memory classes, each class of the second set being a subclass of a class of the first set and being platform-dependent.




FIG. 1

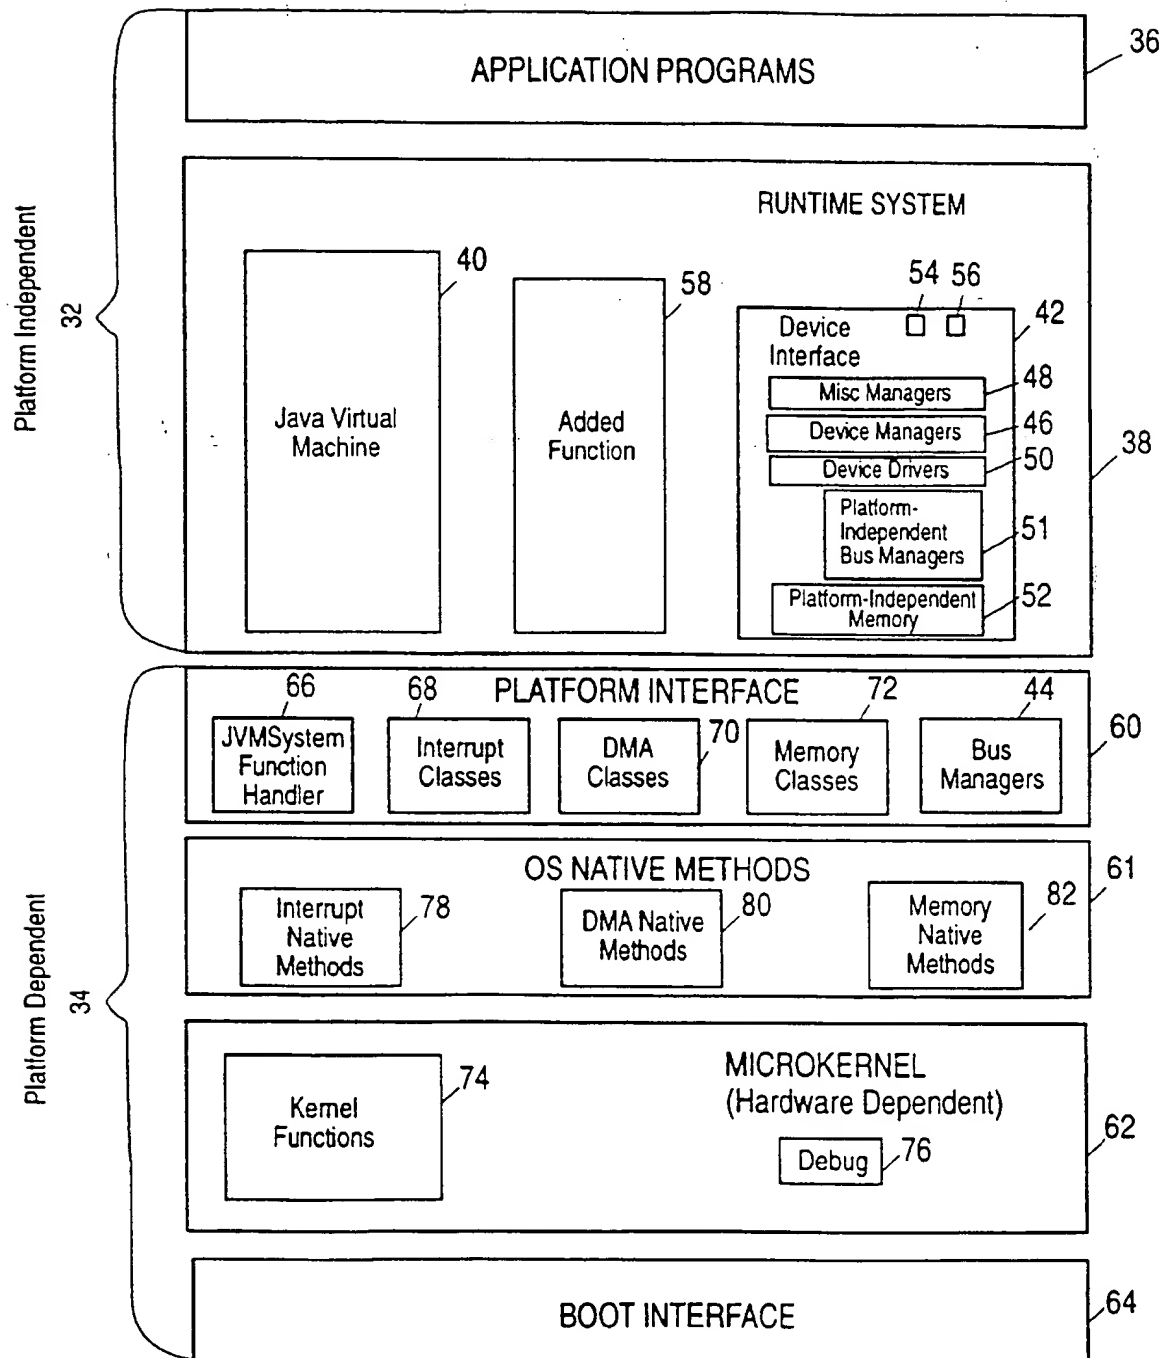


FIG. 2

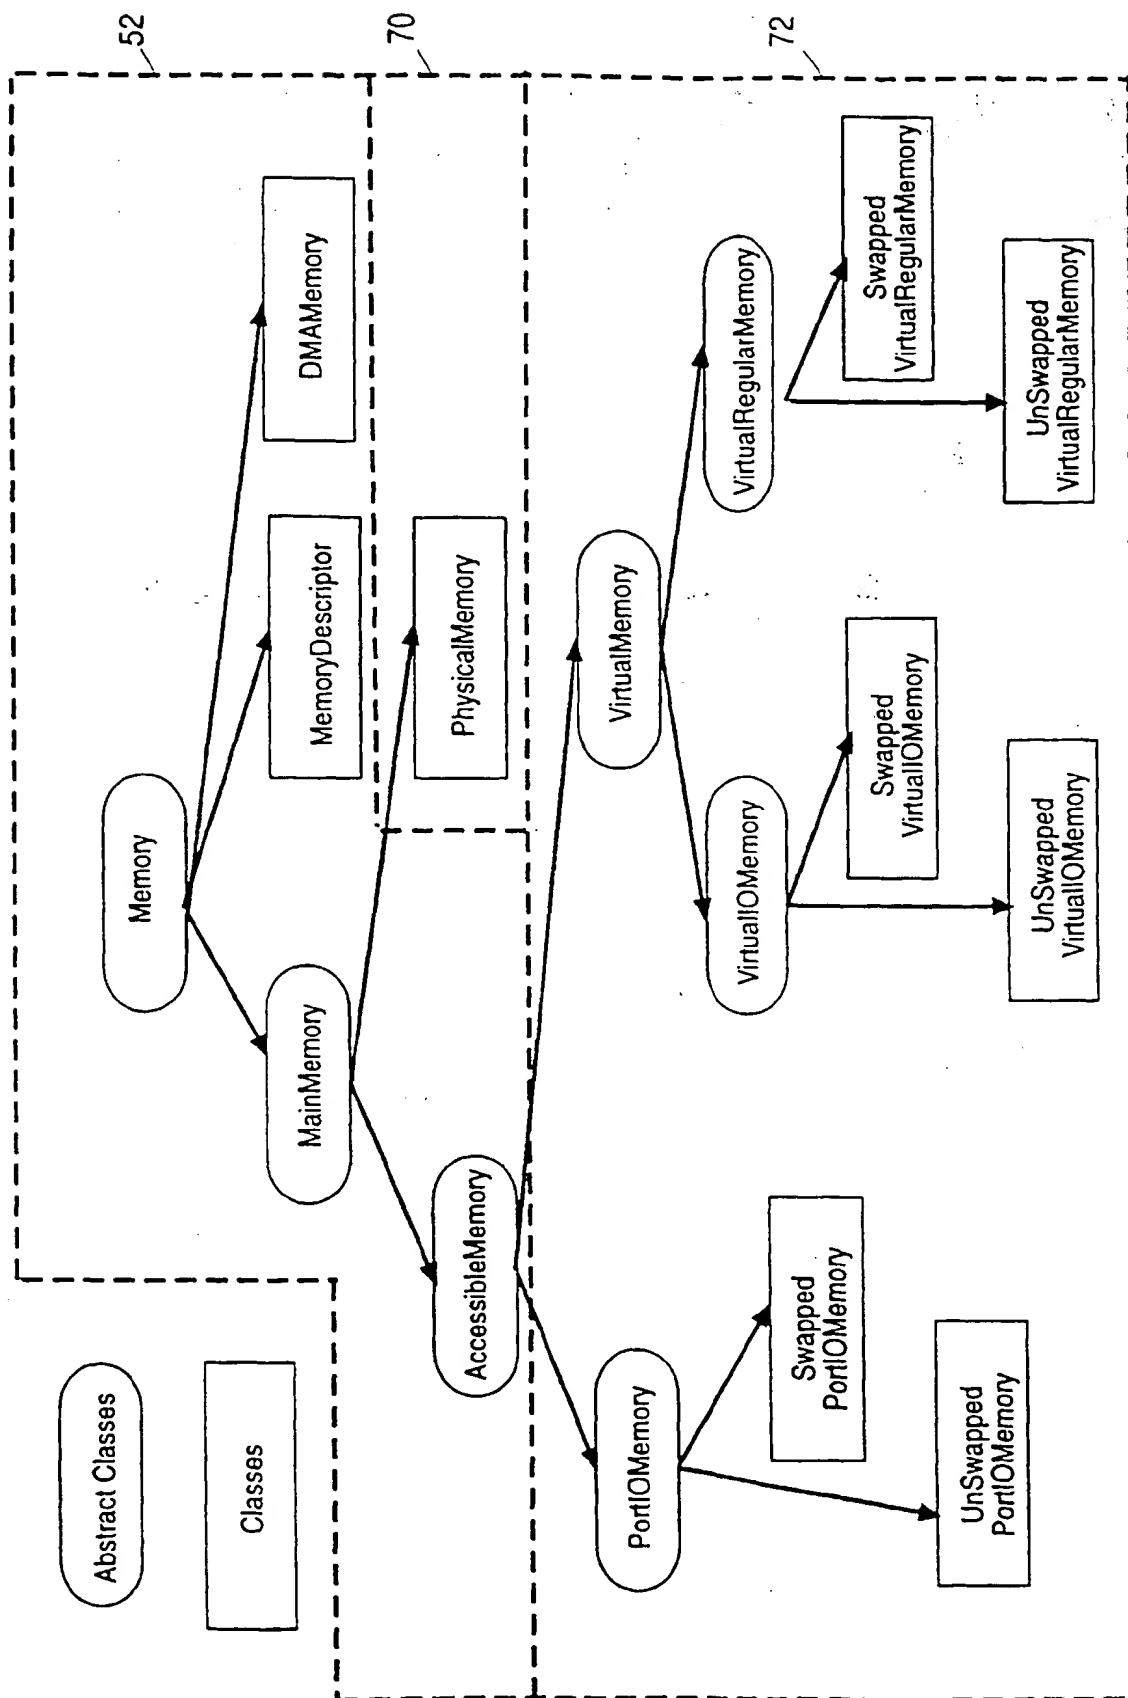


FIG. 3